arXiv:1608.05432v2 [math.AT] 13 Apr 2017

PERSISTENT HOMOLOGY OF ASYMMETRIC NETWORKS: AN APPROACH BASED ON DOWKER FILTRATIONS

SAMIR CHOWDHURY AND FACUNDO MÉMOLI

ABSTRACT. We propose methods for computing two network features with topological underpinnings: the Rips and Dowker Persistent Homology Diagrams. Our formulations work for general networks, which may be asymmetric and may have any real number as an edge weight. We study the sensitivity of Dowker persistence diagrams to intrinsic asymmetry in the data, and investigate the stability properties of both the Dowker and Rips persistence diagrams. We include detailed experiments run on a variety of simulated and real world datasets using our methods.

CONTENTS

1. Introduction 1
1.1. Overview of our approach 3
1.2. Implementations 4
1.3. Organization of the paper 4
1.4. Notation 4
2. Background on persistent homology 4
2.1. Interpolating between $\mathbb{N}$ and $\mathbb{R}$-indexed persistence vector spaces. 6
2.2. Persistence diagrams and barcodes 7
2.3. Interleaving distance and stability of persistence vector spaces. 8
3. Background on networks and our network distance 9
4. The Rips complex of a network 11
5. The Dowker complex of a network 13
5.1. Dowker duality and equivalence of diagrams 15
5.2. Dowker persistence diagrams capture asymmetry 19
6. Rips and Dowker hierarchical clustering methods 29
6.1. The Rips nHCM 31
6.2. The Dowker nHCM 31
7. Implementation and experiments on classification and exploratory data analysis 32
7.1. Simulated hippocampal networks 32
7.2. U.S. economy input-output accounts 34
7.3. U.S. migration 41
7.4. Global migration 47
8. Discussion 49
References 51

1. INTRODUCTION

Networks are used throughout the sciences for representing the complex relations that exist between the objects of a dataset [New03, EK10]. Network data arises from applications in social science [KNT10,

Date: April 17, 2017.
2010 *Mathematics Subject Classification.* 55U99, 68U05 55N35.

EK10], commerce and economy [EGJ14, EK10, AOTS15], neuroscience [Spo11, Spo12, SK04, RS10, Pes14], biology [BO04, HRS10], and defence [Mas14], to name a few sources. Networks are most often directed, in the sense that weights attached to edges do not satisfy any symmetry property, and this asymmetry often precludes the applicability of many standard methods for data analysis.

Network analysis problems come in a wide range of flavors. One problem is in *exploratory data analysis*: given a network representing a dataset of societal, economic, or scientific value, the goal is to obtain insights that are meaningful to the interested party and can help uncover interesting phenomena. Another problem is *network classification*: given a "bag" of networks representing multiple instances of different phenomena, one wants to obtain a clustering which groups the networks together according to the different phenomena they represent.

Because networks are often too complex to deal with directly, one typically extracts certain invariants of networks, and infers structural properties of the networks from properties of these invariants. While there are numerous such network invariants in the existing literature, there is growing interest in adopting a particular invariant arising from *persistent homology* [Fro92, Rob99, ELZ02, ZC05], known as a *persistence diagram*, to the setting of networks. Persistence diagrams are used in the context of finite metric space or point cloud data to pick out *features of significance* while rejecting random noise [Car09]. Since a network on n nodes is regarded, in the most general setting, as an $n \times n$ matrix of real numbers, i.e. as a generalized metric space, it is conceivable that one should be able to describe persistence diagrams for networks as well.

The motivation for computing persistence diagrams of networks is at least two-fold: (1) comparing persistence diagrams has been shown to be a viable method for *shape matching* applications [Fro92, FL99, CZCG04, CSEH07, CZCG05, CCSG$^+$09b], analogous to the network classification problem described above, and (2) persistence diagrams have been successfully applied to feature detection, e.g. in detecting the structure of protein molecules (see [KPT07, XW14] and [ELZ02, §6]) and solid materials (see [HNH$^+$16]) and might thus be a useful tool for exploratory analysis of network datasets.

We point the reader to [Ghr08, EH08, Car09, Wei11] for surveys of persistent homology and its applications.

Some extant approaches that obtain persistence diagrams from networks assume that the underlying network data actually satisfies metric properties [LCK$^+$11, KKC$^+$14]. A more general approach for obtaining persistence diagrams from networks is followed in [HMR09, CH13, GPCI15, PSDV13], albeit with the restriction that the input data sets are required to be symmetric matrices.

Our chief goal is to devise notions of persistent homology that are directly applicable to asymmetric networks in the most general sense, and are furthermore capable of absorbing structural information contained in the asymmetry.

In this paper, we study two types of persistence diagrams: the *Rips* and *Dowker* diagrams. We define both of these invariants in the setting of asymmetric networks with real-valued weights, without assuming any metric properties at all (not symmetry and not even that the matrix vanishes on the diagonal). As a key step in defining the Dowker persistence diagram, we first define two dual constructions, each of which can be referred to as a Dowker persistence diagram, and then prove a result called *Dowker duality* showing that these two possible diagrams are equivalent. Following the line of work in [CCSG$^+$09b], where stability of Rips persistence diagrams arising from finite metric spaces was first established, we formulate similar stability results for the Rips and Dowker persistence diagrams of a network. Through various examples, in particular a family of *cycle networks*, we espouse the idea that 1-dimensional Dowker persistence diagrams are more appropriate than 1-dimensional Rips persistence diagrams for studying asymmetric networks. We note that neither of the 0-dimensional Rips or Dowker persistence diagrams are sufficiently sensitive to asymmetry, and investigate related constructions in the framework of *hierarchical clustering*. We then test these methods on a variety of datasets, and exhibit our approaches for: (1) solving a network classification problem on a database of simulated hippocampal networks, and (2) performing exploratory data analysis on a U.S. economic dataset, a U.S. migration dataset, and a global migration dataset.

1.1. **Overview of our approach.** The first step in constructing a persistence diagram from a network is to construct a nested sequence of simplicial complexes, which, in our work, will be the *Rips* and *Dowker* complexes of a network. Rips and Dowker complexes are classically defined for metric spaces [Ghr14, §2], and the generalization to networks that we use has a simple description. After producing the simplicial complexes, the standard framework of *persistent homology* takes over, and we obtain the Rips or Dowker persistence diagrams.

However, producing these persistence diagrams is not enough. In order for these invariants to be useful in practice, one must verify that the diagrams are *stable* in the following sense: the dissimilarity between two Rips (resp. Dowker) persistence diagrams obtained from two networks should be bounded above by a function of the dissimilarity between the two networks. To our knowledge, stability is not addressed in the existing literature on producing persistence diagrams from networks. In our work, we provide stability results for both the Rips and Dowker persistence diagrams (Propositions 7 and 10). One key ingredient in our proof of this result is a notion of *network distance* that follows previous work in [CMRS14, CM15, CM16a]. This network distance is analogous to the Gromov-Hausdorff distance between metric spaces, which has previously been used to prove stability results for hierarchical clustering [CM08, CM10] and Rips persistence diagrams obtained from finite metric spaces [CCSG$^+$09b, Theorem 3.1]. The Gromov-Hausdorff distance was later used in conjunction with the Algebraic Stability Theorem of [CCSG$^+$09a] to provide alternative proofs of stability results for Rips and Dowker persistence diagrams arising from metric spaces [CDSO14]. Our proofs also use this Algebraic Stability Theorem, but the novelty of our approach lies in a reformulation of the network distance (Proposition 4) that yields direct maps between two networks, thus passing naturally into the machinery of the Algebraic Stability Theorem (without having to define auxiliary constructions such as multivalued maps, as in [CDSO14]).

Practitioners of persistent homology might recall that there are *two* Dowker complexes [Ghr14, p. 73], which we describe as the *source* and *sink* Dowker complexes. A subtle point to note here is that each of these Dowker complexes can be used to construct a persistence diagram. A folklore result in the literature about persistent homology of metric spaces, known as *Dowker duality*, is that the two persistence diagrams arising this way are equal [CDSO14, Remark 4.8]. In this paper we provide a complete proof of this duality in a context strictly more general than that of metric spaces. Our proof imports ideas used by Dowker in his 1952 paper on the homology of a relation [Dow52] and renders them in a persistence framework.

Dowker complexes are also known to researchers who use Q-analysis to study social networks [Joh13, Atk75, Atk72]. We perceive that viewing Dowker complexes through the modern lens of persistence will enrich the classical framework of Q-analysis by incorporating additional information about the *meaningfulness* of features, thus potentially opening new avenues in the social sciences.

A crucial issue that we point out in this paper is that even though we can construct both Rips and Dowker persistence diagrams out of asymmetric networks, the Rips persistence diagrams appear to be *blind* to asymmetry, whereas the Dowker persistence diagrams do exhibit sensitivity to asymmetry. This can be seen "intuitively" from the definitions of the Rips and Dowker complexes. In order to ground this intuition more concretely, we consider a family of highly asymmetric networks, the *cycle networks*, and prove a characterization result for the 1-dimensional Dowker persistence diagram of any network belonging to this family. Some of our experimental results suggest that the Rips persistence diagrams of this family of networks are pathological, in the sense that they do not represent the signatures one would expect from the underlying dataset, which is a directed circle. Dowker persistence diagrams, on the other hand, are well-behaved in this respect in that they succeed at capturing relevant features.

Even though we prove the Dowker duality result soon after introducing Dowker complexes, we note a result later on in the paper that is analogous to the 0-dimensional case of Dowker duality and requires very little machinery to prove. To describe this viewpoint, we segue into the theory of unsupervised learning: we define two *network hierarchical clustering methods*, relying on the Dowker source and sink complexes, respectively, and prove that the output *dendrograms* of these two methods are equivalent. We explain why the dendrogram produced by such a hierarchical clustering method is a stronger network invariant than the 0-dimensional Dowker persistence diagram. One can verify that the 0-dimensional Dowker duality result

follows as a consequence of this more general result about dendrograms, although we do not go into detail about this claim. We also define a Rips network hierarchical clustering method, and remark that the Dowker and Rips network hierarchical clustering methods correspond to the *unilateral* and *reciprocal* clustering methods that are well-studied in the machine learning literature [CMRS13].

An announcement of some of our work will appear in [CM16c].

1.2. **Implementations.** This paper is intended to guide researchers on using persistence diagrams for network analysis, and so we provide details on a variety of implementations that were of interest to us.

The first implementation is in the setting of classifying simulated hippocampal networks, following work in [CI08, DMFC12]. We simulate the activity pattern of hippocampal cells in an animal as it moves around arenas with a number of obstacles, and compile this data into a network which can be interpreted as the transition matrix for the time-reversal of a Markov process. The motivating idea is to ascertain whether the brain can determine, by just looking back at its hippocampal activity and not using any higher reasoning ability, the number of obstacles in the arena the that animal has just finished traversing. The results of computing Dowker persistence diagrams suggest that the hippocampal activity is indeed sufficient to accurately count the number of obstacles in each arena.

Next we consider a network obtained from an input-output account of U.S. economic data. Economists use such data to determine the process by which goods are produced and distributed across various industrial sectors in the U.S. By computing the 1-dimensional Dowker persistence diagram of this network, we are able to obtain asymmetric "flows" of investment across industries.

As another implementation, we consider a network representing U.S. migration. From the 1-dimensional Dowker persistence diagram of this network, we are able to obtain migration flows representing people who do not have a common preferred destination. Quite possibly, this can be attributed to the heterogeneity and diversity of the U.S. For a broader overview of these phenomena, we also study a network representing global migration between 231 administrative regions around the world.

While our analysis of each of these experiments is by no means exhaustive, we perceive that experts in the respective fields will be able to glean more insight from our work. Our datasets and software will be made available on `https://research.math.osu.edu/networks/Datasets.html`.

1.3. **Organization of the paper.** The paper is organized as follows. Notation used globally is defined directly below. §2 contains the necessary background on persistent homology. §3 contains our formulations for networks, as well as some key ingredients of our stability results. §4 contains details about the Rips persistence diagram. The first part of §5 contains details about the Dowker persistence diagram. §5.1 contains the result that we have referred to above as Dowker duality. §5.2.1 contains a family of asymmetric networks, the *cycle networks*, and a full characterization of their 1-dimensional Dowker persistence diagrams. §6 contains details about the Dowker and Rips network hierarchical clustering methods, and an alternative proof of the 0-dimensional Dowker duality result which follows from a stronger result about the Dowker network hierarchical clustering method. Finally, in §7 we provide details on four implementations of our methods to simulated and real-world datasets, as well as some interpretations suggested by our analysis.

1.4. **Notation.** We will write $\mathbb{K}$ to denote a field, which we will fix and use throughout the paper. We will write $\mathbb{Z}_+$ and $\mathbb{R}_+$ to denote the nonnegative integers and reals, respectively. The extended real numbers $\mathbb{R} \cup \{\infty, -\infty\}$ will be denoted $\overline{\mathbb{R}}$. The cardinality of a set X will be denoted $\mathrm{card}(X)$. The collection of nonempty subsets of a set X will be denoted $\mathrm{pow}(X)$. The natural numbers $\{1, 2, 3, \ldots\}$ will be denoted by $\mathbb{N}$. The dimension of a vector space V will be denoted $\dim(V)$. The rank of a linear transformation f will be denoted $\mathrm{rank}(f)$. An isomorphism between vector spaces V and W will be denoted $V \cong W$.

2. Background on persistent homology

Given a finite set X, a *simplicial complex* is defined to be a collection of elements in $\mathrm{pow}(X)$ such that whenever $\sigma \in \mathrm{pow}(X)$ belongs to the collection, any subset $\tau \subseteq \sigma$ belongs to the collection as well. The singleton elements in this collection are the *vertices* of the simplicial complex, the two-element subsets of

X belonging to this collection are the *edges*, and for any $k \in \mathbb{Z}_+$, the $(k+1)$ element subsets of X in this collection are the k-simplices. Whenever we write a k-simplex $\{x_0, x_1, \ldots, x_k\}$, we will assume that the simplex is *oriented* by the ordering $x_0 < x_1 < \ldots < x_k$. We will write $[x_0, x_1, \ldots, x_k]$ to denote the equivalence class of the even permutations of this chosen ordering, and $-[x_0, x_1, \ldots, x_k]$ to denote the equivalence class of the odd permutations of this ordering.

Given two simplicial complexes Σ, Ξ with vertex sets $V(\Sigma), V(\Xi)$, a map $f : V(\Sigma) \to V(\Xi)$ is *simplicial* if $f(\sigma) \in \Xi$ for each $\sigma \in \Sigma$. We will often refer to such a map as a *simplicial map* from Σ to Ξ, denoted as $f : \Sigma \to \Xi$. Note that if $\tau \subseteq \sigma \in \Sigma$, we also have $f(\tau) \subseteq f(\sigma) \in \Xi$.

Let Σ be a finite simplicial complex. For any dimension $k \in \mathbb{Z}_+$, a *k-chain* in Σ is a formal linear combination of oriented k-simplices in Σ, written as $\sum_i a_i \sigma_i$, where each $a_i \in \mathbb{K}$. The collection of all k-chains is a $\mathbb{K}$-vector space (more specifically, a free vector space over $\mathbb{K}$), denoted $C_k(\Sigma)$ or just C_k, and is called the k-chain vector space of Σ. Note that C_k is generated by the (finitely many) k-simplices of Σ. We also define $C_k := \{0\}$ for negative integers k.

Given any k-chain vector space, the *boundary map* $\partial_k : C_k \to C_{k-1}$ is defined as:

$$\partial_k[x_0, \ldots, x_k] = \sum_i (-1)^i [x_0, \ldots, \hat{x}_i, \ldots, x_k], \text{ where } \hat{x}_i \text{ denotes omission of } x_i \text{ from the sequence.}$$

A standard observation here is that $\partial_{k-1} \circ \partial_k = 0$, for any $k \in \mathbb{N}$ [Mun84]. Next, a *chain complex* is defined to be a sequence of vector spaces $\mathcal{C} = (C_k, \partial_k)_{k \in \mathbb{Z}_+}$ with boundary maps such that $\partial_{k-1} \circ \partial_k = 0$. Given a chain complex $\mathcal{C}$ and any $k \in \mathbb{Z}_+$, one may define the following subspaces:

$$Z_k(\mathcal{C}) := \ker(\partial_k) = \{c \in C_k : \partial_k c = 0\}, \text{ the } k\text{-cycles},$$
$$B_k(\mathcal{C}) := \mathrm{im}(\partial_{k+1}) = \{c \in C_k : c = \partial_{k+1} b\}, \text{the } k\text{-boundaries}.$$

The quotient vector space $H_k(\mathcal{C}) := Z_k(\mathcal{C})/B_k(\mathcal{C})$ is called the *k-th homology of the chain complex $\mathcal{C}$*. The dimension of $H_k(\mathcal{C})$ is called the *k-th Betti number* of $\mathcal{C}$, denoted $\beta_k(\mathcal{C})$.

Given two chain complexes $\mathcal{C} = (C_k, \partial_k^{\mathcal{C}})_{k \in \mathbb{Z}_+}$ and $\mathcal{C}' = (C'_k, \partial^{\mathcal{C}'})_{k \in \mathbb{Z}_+}$, a *chain map* $\Phi : \mathcal{C} \to \mathcal{C}'$ is a sequence of linear maps $\{\varphi_k : C_k \to C'_k\}_{k \in \mathbb{Z}_+}$ such that $\partial_k^{\mathcal{C}'} \circ \varphi_k = \varphi_{k-1} \circ \partial_k^{\mathcal{C}}$ for each $k \in \mathbb{Z}_+$. By virtue of this property, φ_k maps $Z_k(\mathcal{C})$ to $Z_k(\mathcal{C}')$ and $B_k(\mathcal{C})$ to $B_k(\mathcal{C}')$ for each $k \in \mathbb{Z}_+$. Thus a chain map Φ induces a natural sequence of maps $\Phi_\# = \{(\varphi_k)_\# : H_k(\mathcal{C}) \to H_k(\mathcal{C}')\}_{k \in \mathbb{Z}_+}$ between homology vector spaces. Finally, we note that simplicial maps between simplicial complexes induce natural chain maps between the corresponding chain complexes [Mun84, §1.12]. More specifically, given two simplicial complexes Σ, Ξ and a simplicial map $f : \Sigma \to \Xi$, there exists a natural chain map $f_* : \mathcal{C}(\Sigma) \to \mathcal{C}(\Xi)$, which in turn induces a linear map $(f_k)_\# : H_k(\mathcal{C}(\Sigma)) \to H_k(\mathcal{C}(\Xi))$ for each $k \in \mathbb{Z}_+$.

There are two special properties of the chain maps induced by simplicial maps that we will use throughout the paper. These properties are often referred to as *functoriality of homology*. Let Σ be a simplicial complex, and let $\iota : \Sigma \to \Sigma$ denote the identity simplicial map. Let $\{(\iota_k)_* : C_k(\Sigma) \to C_k(\Sigma)\}_{k \in \mathbb{Z}_+}$ denote the chain map induced by ι. Then for each $k \in \mathbb{Z}_+$, $(\iota_k)_\# : H_k(\mathcal{C}(\Sigma)) \to H_k(\mathcal{C}(\Sigma))$ is the identity map [Mun84, Theorem 12.2]. Next let Ξ, Θ be two more simplicial complexes, and let $f : \Sigma \to \Xi, g : \Xi \to \Theta$ be two simplicial maps. Let $\{(f_k)_* : C_k(\Sigma) \to C_k(\Xi)\}_{k \in \mathbb{Z}_+}$ and $\{(g_k)_* : C_k(\Xi) \to C_k(\Theta)\}_{k \in \mathbb{Z}_+}$ be the chain maps induced by f, g. Then we have [Mun84, Theorem 12.2]:

$$(g_k \circ f_k)_\# = (g_k)_\# \circ (f_k)_\# \qquad \text{for each } k \in \mathbb{Z}_+. \tag{1}$$

The operations we have described above, i.e. that of passing from simplicial complexes and simplicial maps to chain complexes and induced chain maps, and then to homology vector spaces with induced linear maps, will be referred to as *passing to homology*.

To introduce the idea of persistence, let $(\delta_i)_{i \in \mathbb{N}}$ be an increasing sequence of real numbers. A *filtration* of Σ (also called a *filtered simplicial complex*) is defined to be an increasing sequence $(\Sigma^{\delta_i})_{i \in \mathbb{N}}$ of simplicial complexes, such that:

$$\Sigma^{\delta_i} \subseteq \Sigma^{\delta_{i+1}} \text{ for each } i \in \mathbb{N}, \quad \Sigma^{\delta_1} = \varnothing, \quad \Sigma^{\delta_n} = \Sigma \text{ for some } n \in \mathbb{N}, \text{ and } \Sigma^{\delta_i} = \Sigma \text{ for each } i \geqslant n.$$

Next, given an increasing sequence of real numbers $(\delta_i)_{i\in\mathbb{N}}$, a *persistence vector space* is defined to be a family of vector spaces $\mathcal{V} = \{V^{\delta_i} \xrightarrow{\nu_{i,i+1}} V^{\delta_{i+1}}\}_{i\in\mathbb{N}}$ with linear maps between them, such that: (1) $\dim(V^{\delta_i}) < \infty$ for each $i \in \mathbb{N}$, and (2) for each $i \geqslant n$ for some $n \in \mathbb{N}$, we have vector space isomorphisms $V^{\delta_i} \cong V^{\delta_{i+1}}$. Each δ_i is referred to as a *resolution* parameter. Since the family $\mathcal{V}$ is indexed by a countable sequence of resolution parameters, we denote the collection of all such persistence vector spaces by $\mathbf{PVec}(\mathbb{N})$.

Note that at each step Σ^{δ_i} of the filtered simplicial complex $(\Sigma^{\delta_i})_{i\in\mathbb{N}}$ above, we may produce an associated chain complex $\mathcal{C}(\Sigma^{\delta_i})$. The inclusion maps $\iota_{i,i+1} : \Sigma^{\delta_i} \hookrightarrow \Sigma^{\delta_{i+1}}$ at the simplicial level become linear transformations between vector spaces at the chain complex level. Thus we obtain a family of chain complexes $\{\mathcal{C}^{\delta_i} \xrightarrow{(\iota_{i,i+1})_*} \mathcal{C}^{\delta_{i+1}}\}_{i\in\mathbb{N}}$ with linear maps between them. Then by taking the kth homology vector space of this family, for a given dimension $k \in \mathbb{Z}_+$, we obtain the kth persistence vector space associated to $(\Sigma^{\delta_i})_{i\in\mathbb{N}}$, denoted

$$\mathcal{H}_k(\Sigma) := \{H_k(\mathcal{C}^{\delta_i}) \xrightarrow{(\iota_{i,i+1})_\#} H_k(\mathcal{C}^{\delta_{i+1}})\}_{i\in\mathbb{N}}.$$

One can use the finiteness of Σ to check that the family of vector spaces $\mathcal{H}_k(\Sigma)$ satisfies the conditions defining a persistence vector space.

So far, we have defined persistence vector spaces that are indexed by a countable sequence of resolution parameters, i.e. $\mathbb{N}$-indexed persistence vector spaces. However, certain results in the persistent homology that we will use throughout the paper are stated for $\mathbb{R}$-indexed persistence vector spaces. We will now define these objects, and show how to interpolate between these two notions.

2.1. **Interpolating between $\mathbb{N}$ and $\mathbb{R}$-indexed persistence vector spaces.** Let $\mathcal{V} = \{V^{\delta_i} \xrightarrow{\nu_{i,i+1}} V^{\delta_{i+1}}\}_{i\in\mathbb{N}} \in \mathbf{PVec}(\mathbb{N})$. We can now define a family of vector spaces indexed by $\mathbb{R}$ as follows:

$$U^\delta := V^{\delta_i} \text{ whenever } \delta \in [\delta_i, \delta_{i+1}) \text{ for some } i \in \mathbb{N}.$$

For convenience, define the following map for any $j \geqslant i+1 > i \in \mathbb{N}$:

$$\nu_i^j := \nu_{j-1,j} \circ \nu_{j-2,j-1} \circ \cdots \circ \nu_{i,i+1}.$$

Note ν_i^j is a linear map from V^{δ_i} to V^{δ_j}. Then for any $\delta \leqslant \delta'$, we can define a linear map $\mu_{\delta,\delta'} : U^\delta \to U^{\delta'}$ as follows:

$$\mu_{\delta,\delta'} := \begin{cases} \mathrm{id} : V^{\delta_i} \to V^{\delta_i} & : \delta, \delta' \in [\delta_i, \delta_{i+1}) \text{ for some } i \in \mathbb{N}, \\ \nu_i^j : V^{\delta_i} \to V^{\delta_j} & : \delta \in [\delta_i, \delta_{i+1}), \delta' \in [\delta_j, \delta_{j+1}) \text{ for some } i < j \in \mathbb{N}. \end{cases}$$

Note that: (1) $\dim(U^\delta) < \infty$ at each $\delta \in \mathbb{R}$, (2) all maps $\mu_{\delta,\delta'}$ are isomorphisms for sufficiently large δ, δ', and (3) there are only finitely many values of $\delta \in \mathbb{R}$ such that $U^{\delta-\varepsilon} \ncong U^\delta$ for each $\varepsilon > 0$. An $\mathbb{R}$-indexed family of vector spaces with linear maps satisfying these three conditions is called an $\mathbb{R}$-*indexed persistence vector space*. The collection of all such families is denoted $\mathbf{PVec}(\mathbb{R})$.

Thus far, we have defined a method $\Phi : \mathbf{PVec}(\mathbb{N}) \to \mathbf{PVec}(\mathbb{R})$ of passing from $\mathbf{PVec}(\mathbb{N})$ to $\mathbf{PVec}(\mathbb{R})$. We now reverse this construction. Let $\{U^\delta \xrightarrow{\mu_{\delta,\delta'}} U^{\delta'}\}_{\delta\leqslant\delta'} \in \mathbf{PVec}(\mathbb{R})$. Let $\{\delta_1, \delta_2, \ldots, \delta_n\}$ be the finite set of resolution parameters at which U^δ undergoes a change. Then for each $i \in \mathbb{N}$, we can define:

$$V^{\delta_i} := \begin{cases} U^{\delta_i} & : 1 \leqslant i \leqslant n \\ U^{\delta_n} & : i > n. \end{cases}$$

We also define $\nu_{i,i+i} := \mu_{\delta_i,\delta_{i+1}}$ for each $i \in \mathbb{N}$. This yields an element of $\mathbf{PVec}(\mathbb{N})$, as desired. Thus we have a method $\Psi : \mathbf{PVec}(\mathbb{R}) \to \mathbf{PVec}(\mathbb{N})$ of passing from $\mathbf{PVec}(\mathbb{R})$ to $\mathbf{PVec}(\mathbb{N})$.

Note that the elements in both $\mathbf{PVec}(\mathbb{N})$ and $\mathbf{PVec}(\mathbb{R})$ involve only a finite number of vector spaces, up to isomorphism. By our construction of Φ and Ψ, and the classification results in [CZCG05, §5.2], it follows that Φ and Ψ preserve a certain invariant, called a *barcode*, of each element of $\mathbf{PVec}(\mathbb{N})$ and $\mathbf{PVec}(\mathbb{R})$, respectively. Thus when proving results about barcodes, one can define constructions on $\mathbf{PVec}(\mathbb{R})$ and

transport them to $\mathbf{PVec}(\mathbb{N})$, and vice versa. In the following section, we go into detail about the barcodes referred to above.

2.2. **Persistence diagrams and barcodes.** To each persistence vector space, one may associate a multiset of intervals, called a *persistence barcode* or *persistence diagram*. This barcode is a full invariant of a persistence vector space [ZC05], and it has the following natural interpretation: given a barcode corresponding to a persistence vector space obtained from a filtered simplicial complex Σ, the long bars correspond to meaningful features of Σ, whereas the short bars correspond to noise or artifacts in the data. The standard treatment of persistence barcodes and diagrams appear in [ZC05] and [ELZ02]. We follow a more modern presentation that appeared in [EJM15]. To build intuition, we refer the reader to Figure 1.

Let $\mathcal{V} = \{V^{\delta_i} \xrightarrow{\nu_{i,i+1}} V^{\delta_{i+1}})_{i\in\mathbb{N}} \in \mathbf{PVec}(\mathbb{N})$. Because all but finitely many of the $\nu_{\bullet,\bullet+1}$ maps are isomorphisms, one may choose a basis $(B_i)_{i\in\mathbb{N}}$ for each V^{δ_i}, $i \in \mathbb{N}$, such that $\nu_{i,i+1}|_{B_i}$ is injective for each $i \in \mathbb{N}$, and

$$\operatorname{rank}(\nu_{i,i+1}) = \operatorname{card}(\operatorname{im}(\nu_{i,i+1}|_{B_i}) \cap B_{i+1}), \text{ for each } i \in \mathbb{N} \text{ [EJM15, Basis Lemma].}$$

Here $\nu_{i,i+1}|_{B_i}$ denotes the restriction of $\nu_{i,i+1}$ to the set B_i. Fix such a collection $(B_i)_{i\in\mathbb{N}}$ of bases. Next define:

$$L := \{(b,i) : b \in B_i, b \notin \operatorname{im}(\nu_{i-1,i}), i \in \{2,3,4,\ldots\}\} \cup \{(b,1) : b \in B_1\}.$$

Next define a map $\ell : L \to \mathbb{N}$ as follows:

$$\ell(b,i) := \sup\{k \in \mathbb{N} : \nu_i^k(b) \in B_{i+1}, b \in B_i\}.$$

The *persistence barcode of* $\mathcal{V}$ is then defined to be the following multiset of intervals

$$\mathbf{Pers}(\mathcal{V}) := \big[[\delta_i, \delta_{j+1}) : \text{there exists} (b,i) \in L \text{ such that } \ell(b,i) = j\big],$$

where the bracket notation denotes taking the multiset and the multiplicity of $[\delta_i, \delta_{j+1})$ is the number of elements $(b,i) \in L$ such that $\ell(b,i) = j$.

These intervals, which are called *persistence intervals*, are then represented as a set of lines over a single axis. Equivalently, the intervals in $\mathbf{Pers}(\mathcal{V})$ can be visualized as a multiset of points lying on or above the diagonal in $\overline{\mathbb{R}}^2$, counted with multiplicity. This is the case for the *persistence diagram of* $\mathcal{V}$, which is defined as follows:

$$\operatorname{Dgm}(\mathcal{V}) := \big[(\delta_i, \delta_{j+1}) \in \overline{\mathbb{R}}^2 : [\delta_i, \delta_{j+1}) \in \mathbf{Pers}(\mathcal{V})\big],$$

where the multiplicity of $(\delta_i, \delta_{j+1}) \in \overline{\mathbb{R}}^2$ is given by the multiplicity of $[\delta_i, \delta_{j+1}) \in \mathbf{Pers}(\mathcal{V})$.

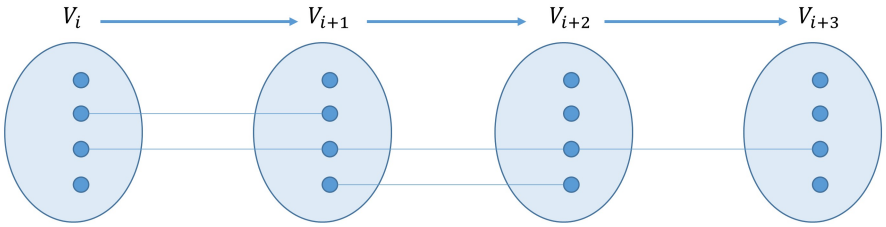

FIGURE 1. Intuition behind a persistence barcode. Let $i \in \mathbb{N}$, and consider a sequence of vector spaces $V_i, V_{i+1}, V_{i+2}, V_{i+3}$ as above, with linear maps $\{\nu_{i,i+1}, \nu_{i+1,i+2}, \nu_{i+2,i+3}\}$. The dark dots represent basis elements, where the bases are chosen such that $\nu_{i,i+1}$ maps the basis elements of V_i to those of V_{i+1}, and so on. Such a choice of basis is possible by performing row and column operations on the matrices of the linear maps [EJM15, Basis Lemma]. The persistence barcode of this sequence can then be read off from the "strings" joining the dots. In this case, the barcode is the collection $\{[i, i+1], [i, i+3], [i+1, i+2]\}$. Note that when these intervals are read in $\mathbb{Z}$, they are the same as the half-open intervals one would expect from the definition of the persistence barcode given above.

The *bottleneck distance* between persistence diagrams, and more generally between multisets A, B of points in $\overline{\mathbb{R}}^2$, is defined as follows:

$$d_{\mathrm{B}}(A,B) := \inf \left\{ \sup_{a\in A} \|a - \varphi(a)\|_\infty : \varphi : A \cup \Delta^\infty \to B \cup \Delta^\infty \text{ a bijection} \right\}.$$

Here $\|(p,q)-(p',q')\|_\infty := \max(|p-p'|, |q-q'|)$ for each $p, q, p', q' \in \mathbb{R}$, and Δ^∞ is the multiset consisting of each point on the diagonal, taken with infinite multiplicity.

Remark 1. From the definition of bottleneck distance, it follows that points in a persistence diagram $\mathrm{Dgm}(\mathcal{V})$ that belong to the diagonal do not contribute to the bottleneck distance between $\mathrm{Dgm}(\mathcal{V})$ and another diagram $\mathrm{Dgm}(\mathcal{U})$. Thus whenever we describe a persistence diagram as being *trivial*, we mean that either it is empty, or it does not have any off-diagonal points.

There are numerous equivalent ways of formulating the definitions we have provided in this section. For more details, we refer the reader to [CDS10, EJM15, EH10, ELZ02, ZC05, BL14].

2.3. **Interleaving distance and stability of persistence vector spaces.** In what follows, we will consider $\mathbb{R}$-indexed persistence vector spaces $\mathbf{PVec}(\mathbb{R})$.

Given $\varepsilon \geqslant 0$, two $\mathbb{R}$-indexed persistence vector spaces $\mathcal{V} = \{V^\delta \xrightarrow{\nu_{\delta,\delta'}} V^{\delta'}\}_{\delta\leqslant\delta'}$ and $\mathcal{U} = \{U^\delta \xrightarrow{\mu_{\delta,\delta'}} U^{\delta'}\}_{\delta\leqslant\delta'}$ are said to be *ε-interleaved* [CCSG+09a, BL14] if there exist two families of linear maps

$$\{\varphi_{\delta,\delta+\varepsilon} : V^\delta \to V^{\delta+\varepsilon}\}_{\delta\in\mathbb{R}},$$
$$\{\psi_{\delta,\delta+\varepsilon} : U^\delta \to U^{\delta+\varepsilon}\}_{\delta\in\mathbb{R}}$$

such that the following diagrams commute for all $\delta' \geqslant \delta \in \mathbb{R}$:

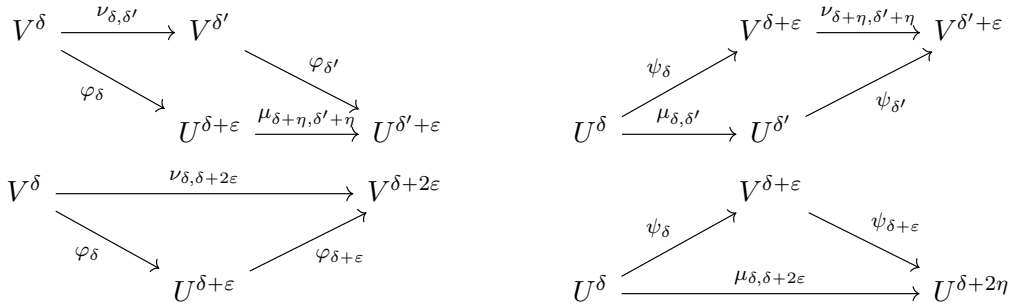

The purpose of introducing ε-interleavings is to define a pseudometric on the collection of persistence vector spaces. The *interleaving distance* between two $\mathbb{R}$-indexed persistence vector spaces $\mathcal{V}, \mathcal{U}$ is given by:

$$d_{\mathrm{I}}(\mathcal{U},\mathcal{V}) := \inf \{\varepsilon \geqslant 0 : \mathcal{U} \text{ and } \mathcal{V} \text{ are } \varepsilon\text{-interleaved}\}.$$

One can verify that this definition induces a pseudometric on the collection of persistence vector spaces [CCSG+09a, BL14]. The interleaving distance can then be related to the bottleneck distance as follows:

Theorem 2 (Algebraic Stability Theorem, [CCSG+09a])**.** *Let $\mathcal{U}, \mathcal{V}$ be two $\mathbb{R}$-indexed persistence vector spaces. Then,*

$$d_{\mathrm{B}}(\mathrm{Dgm}(\mathcal{U}), \mathrm{Dgm}(\mathcal{V})) \leqslant d_{\mathrm{I}}(\mathcal{U},\mathcal{V}).$$

Stability results are at the core of persistent homology, beginning with the classical bottleneck stability result in [CSEH07]. One of our key contributions is to use the Algebraic Stability Theorem stated above, along with Lemma §3 stated below, to prove stability results for methods of computing persistent homology of a network.

Before stating the following lemma, recall that two simplicial maps $f, g : \Sigma \to \Xi$ are *contiguous* if for any simplex $\sigma \in \Sigma$, $f(\sigma) \cup g(\sigma)$ is a simplex of Ξ. Note that if f, g are contiguous maps, then their induced chain maps are chain homotopic, and as a result, the induced maps $f_\#$ and $g_\#$ for homology are equal [Mun84, Theorem 12.5].

Lemma 3 (Stability Lemma)**.** *Let $\mathfrak{F}, \mathfrak{G}$ be two filtered simplicial complexes written as*

$$\left\{\mathfrak{F}^{\delta} \xrightarrow{s_{\delta,\delta'}} \mathfrak{F}^{\delta'}\right\}_{\delta' \geqslant \delta \in \mathbb{R}} \text{ and } \left\{\mathfrak{G}^{\delta} \xrightarrow{t_{\delta,\delta'}} \mathfrak{G}^{\delta'}\right\}_{\delta' \geqslant \delta \in \mathbb{R}},$$

where $s_{\delta,\delta'}$ and $t_{\delta,\delta'}$ denote the natural inclusion maps. Suppose $\eta \geqslant 0$ is such that there exist families of simplicial maps $\left\{\varphi_\delta : \mathfrak{F}^{\delta} \to \mathfrak{G}^{\delta+\eta}\right\}_{\delta \in \mathbb{R}}$ and $\left\{\psi_\delta : \mathfrak{G}^{\delta} \to \mathfrak{F}^{\delta+\eta}\right\}_{\delta \in \mathbb{R}}$ such that the following are satisfied for any $\delta' \geqslant \delta$:

(1) *$t_{\delta+\eta,\delta'+\eta} \circ \varphi_\delta$ and $\varphi_{\delta'} \circ s_{\delta,\delta'}$ are contiguous*
(2) *$s_{\delta+\eta,\delta'+\eta} \circ \psi_\delta$ and $\psi_{\delta'} \circ t_{\delta,\delta'}$ are contiguous*
(3) *$\psi_{\delta+\eta} \circ \varphi_\delta$ and $s_{\delta,\delta+2\eta}$ are contiguous*
(4) *$\varphi_{\delta+\eta} \circ \psi_\delta$ and $t_{\delta,\delta+2\eta}$ are contiguous.*

All the diagrams are as below:

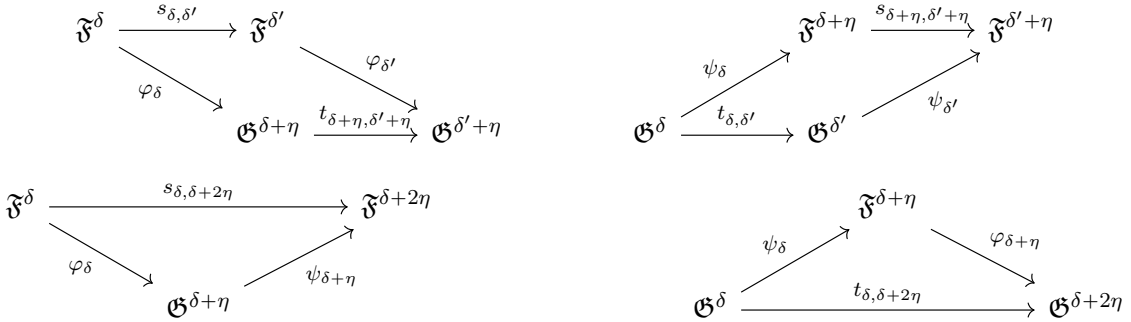

For each $k \in \mathbb{Z}_+$, let $\mathcal{H}_k(\mathfrak{F}), \mathcal{H}_k(\mathfrak{G})$ denote the k-dimensional persistence vector spaces associated to $\mathfrak{F}$ and $\mathfrak{G}$. Then for each $k \in \mathbb{Z}_+$,

$$d_{\mathrm{B}}(\mathrm{Dgm}_k(\mathcal{H}_k(\mathfrak{F})), \mathrm{Dgm}_k(\mathcal{H}_k(\mathfrak{G}))) \leqslant d_{\mathrm{I}}(\mathcal{H}_k(\mathfrak{F}), \mathcal{H}_k(\mathfrak{G})) \leqslant \eta.$$

Proof of Lemma 3. The first inequality holds by the Algebraic Stability Theorem. For the second inequality, note that the contiguous simplicial maps in the diagrams above induce chain maps between the corresponding chain complexes, and these in turn induce equal linear maps at the level of homology vector spaces. To be more precise, first consider the maps $t_{\delta+\eta,\delta'+\eta} \circ \varphi_\delta$ and $\varphi_{\delta'} \circ s_{\delta,\delta'}$. These simplicial maps induce linear maps $(t_{\delta+\eta,\delta'+\eta} \circ \varphi_\delta)_\#, (\varphi_{\delta'} \circ s_{\delta,\delta'})_\# : H_k(\mathfrak{F}^{\delta}) \to H_k(\mathfrak{G}^{\delta'+\eta})$. Because the simplicial maps are assumed to be contiguous, we have:

$$(t_{\delta+\eta,\delta'+\eta} \circ \varphi_\delta)_\# = (\varphi_{\delta'} \circ s_{\delta,\delta'})_\#.$$

By invoking functoriality of homology, we then have:

$$(t_{\delta+\eta,\delta'+\eta})_\# \circ (\varphi_\delta)_\# = (\varphi_{\delta'})_\# \circ (s_{\delta,\delta'})_\#.$$

Analogous results hold for the other pairs of contiguous maps. Thus we obtain commutative diagrams upon passing to homology, and so $\mathcal{H}_k(\mathfrak{F}), \mathcal{H}_k(\mathfrak{G})$ are η-interleaved for each $k \in \mathbb{Z}_+$. Thus we get:

$$d_{\mathrm{I}}(\mathcal{H}_k(\mathfrak{F}), \mathcal{H}_k(\mathfrak{G})) \leqslant \eta.$$

□

3. Background on networks and our network distance

A *network* is a pair (X, ω_X) where X is a finite set and $\omega_X : X \times X \to \mathbb{R}$ is a *weight function*. Note that ω_X need not satisfy the triangle inequality, any symmetry condition, or even the requirement that $\omega_X(x, x) = 0$ for all $x \in X$. The weights are even allowed to be negative. The collection of all such networks is denoted $\mathcal{N}$.

When comparing networks of the same size, e.g. two networks $(X, \omega_X), (X, \omega'_X)$, a natural method is to consider the ℓ^∞ distance:

$$\|\omega_X - \omega'_X\|_\infty := \max_{x,x' \in X} |\omega_X(x, x') - \omega'_X(x, x')|.$$

But one would naturally want a generalization of the ℓ^∞ distance that works for networks having different sizes. In this case, one needs a way to correlate points in one network with points in the other. To see how this can be done, let $(X,\omega_X),(Y,\omega_Y)\in\mathcal{N}$. Let R be any nonempty relation between X and Y, i.e. a nonempty subset of $X\times Y$. The *distortion* of the relation R is given by:

$$\operatorname{dis}(R):=\max_{(x,y),(x',y')\in R}|\omega_X(x,x')-\omega_Y(y,y')|.$$

A *correspondence between X and Y* is a relation R between X and Y such that $\pi_X(R)=X$ and $\pi_Y(R)=Y$, where $\pi_X:X\times Y\to X$ and $\pi_Y:X\times Y\to Y$ denote the natural projections. The collection of all correspondences between X and Y will be denoted $\mathscr{R}(X,Y)$.

Following previous work in [CMRS14, CM15, CM16a] the *network distance* $d_{\mathcal{N}}:\mathcal{N}\times\mathcal{N}\to\mathbb{R}_+$ is then defined as:

$$d_{\mathcal{N}}(X,Y):=\frac{1}{2}\min_{R\in\mathscr{R}}\operatorname{dis}(R).$$

It can be verified that $d_{\mathcal{N}}$ as defined above is a pseudometric, and that the networks at 0-distance can be completely characterized [CM15]. Next we wish to prove the reformulation in Proposition 4. First we define the distortion of a map between two networks. Given any $(X,\omega_X),(Y,\omega_Y)\in\mathcal{N}$ and a map $\varphi:(X,\omega_X)\to(Y,\omega_Y)$, the *distortion* of φ is defined as:

$$\operatorname{dis}(\varphi):=\max_{x,x'\in X}|\omega_X(x,x')-\omega_Y(\varphi(x),\varphi(x'))|.$$

Next, given maps $\varphi:(X,\omega_X)\to(Y,\omega_Y)$ and $\psi:(Y,\omega_Y)\to(X,\omega_X)$, we define two *co-distortion* terms:

$$C_{X,Y}(\varphi,\psi):=\max_{(x,y)\in X\times Y}|\omega_X(x,\psi(y))-\omega_Y(\varphi(x),y)|,$$
$$C_{Y,X}(\psi,\varphi):=\max_{(y,x)\in Y\times X}|\omega_Y(y,\varphi(x))-\omega_X(\psi(y),x)|.$$

Proposition 4. *Let $(X,\omega_X),(Y,\omega_Y)\in\mathcal{N}$. Then,*

$$d_{\mathcal{N}}(X,Y)=\tfrac{1}{2}\inf\{\max(\operatorname{dis}(\varphi),\operatorname{dis}(\psi),C_{X,Y}(\varphi,\psi),C_{Y,X}(\psi,\varphi)):\varphi:X\to Y,\psi:Y\to X\ \textit{any maps}\}.$$

Proof of Proposition 4. First we show that:

$$d_{\mathcal{N}}(X,Y)\geqslant\tfrac{1}{2}\inf\{\max(\operatorname{dis}(\varphi),\operatorname{dis}(\psi),C_{X,Y}(\varphi,\psi),C_{Y,X}(\psi,\varphi)):\varphi:X\to Y,\psi:Y\to X\ \text{any maps}\}.$$

Let $\eta=d_{\mathcal{N}}(X,Y)$, and let R be a correspondence such that $\operatorname{dis}(R)=2\eta$. We can define maps $\varphi:X\to Y$ and $\psi:Y\to X$ as follows: for each $x\in X$, set $\varphi(x)=y$ for some y such that $(x,y)\in R$. Similarly, for each $y\in Y$, set $\psi(y)=x$ for some x such that $(x,y)\in R$. Thus for any $x\in X,y\in Y$, we obtain $|\omega_X(x,\psi(y))-\omega_Y(\varphi(x),y)|\leqslant 2\eta$ and $|\omega_X(\psi(y),x)-\omega_Y(y,\varphi(x))|\leqslant 2\eta$. So we have both $C_{X,Y}(\varphi,\psi)\leqslant 2\eta$ and $C_{Y,X}(\psi,\varphi)\leqslant 2\eta$. Also for any $x,x'\in X$, we have $(x,\varphi(x)),(x',\varphi(x'))\in R$. Thus we also have

$$|\omega_X(x,x')-\omega_Y(\varphi(x),\varphi(x'))|\leqslant 2\eta.$$

So $\operatorname{dis}(\varphi)\leqslant 2\eta$ and similarly $\operatorname{dis}(\psi)\leqslant 2\eta$. This proves the "$\geqslant$" case.

Next we wish to show:

$$d_{\mathcal{N}}(X,Y)\leqslant\tfrac{1}{2}\inf\{\max(\operatorname{dis}(\varphi),\operatorname{dis}(\psi),C_{X,Y}(\varphi,\psi),C_{Y,X}(\psi,\varphi)):\varphi:X\to Y,\psi:Y\to X\ \text{any maps}\}.$$

Suppose φ,ψ are given, and $\frac{1}{2}\max(\operatorname{dis}(\varphi),\operatorname{dis}(\psi),C_{X,Y}(\varphi,\psi),C_{Y,X}(\psi,\varphi))=\eta$.

Let $R_X=\{(x,\varphi(x):x\in X\}$ and let $R_Y=\{(\psi(y),y):y\in Y\}$. Then $R=R_X\cup R_Y$ is a correspondence. We wish to show that for any $z=(a,b),z'=(a',b')\in R$,

$$|\omega_X(a,a')-\omega_Y(b,b')|\leqslant 2\eta.$$

This will show that $\operatorname{dis}(R)\leqslant 2\eta$, and so $d_{\mathcal{N}}(X,Y)\leqslant\eta$.

To see this, let $z,z'\in R$. Note that there are four cases: (1) $z,z'\in R_X$, (2) $z,z'\in R_Y$, (3) $z\in R_X,z'\in R_Y$, and (4) $z\in R_Y,z'\in R_X$. In the first two cases, the desired inequality follows because $\operatorname{dis}(\varphi),\operatorname{dis}(\psi)\leqslant 2\eta$. The inequality follows in cases (3) and (4) because $C_{X,Y}(\varphi,\psi)\leqslant 2\eta$ and $C_{Y,X}(\psi,\varphi)\leqslant 2\eta$, respectively. Thus $d_{\mathcal{N}}(X,Y)\leqslant\eta$. □

Remark 5. Proposition 4 is analogous to a result of Kalton and Ostrovskii [KO97, Theorem 2.1] where—instead of $d_{\mathcal{N}}$—one has the Gromov-Hausdorff distance between metric spaces. We remark that when restricted to the special case of networks that are also metric spaces, the network distance $d_{\mathcal{N}}$ agrees with the Gromov-Hausdorff distance. Details on the Gromov-Hausdorff distance can be found in [BBI01].

An important remark is that in the Kalton-Ostrovskii formulation, there is only one co-distortion term. When Proposition 4 is applied to metric spaces, the two co-distortion terms become equal by symmetry, and thus the Kalton-Ostrovskii formulation is recovered. But *a priori*, the lack of symmetry in the network setting requires us to consider both terms.

Remark 6. In the following sections, we will propose methods for computing persistent homology of networks, and prove that they are stable using Lemma 3. Note that similar results, valid in the setting of metric spaces, have appeared in [CCSG+09b, CDSO14]. However, the proofs in [CDSO14] invoke an auxiliary construction of multivalued maps arising from correspondences, whereas our proofs simply use the maps φ, ψ arising directly from the reformulation of $d_{\mathcal{N}}$ (Proposition 4), thus streamlining the treatment.

We are now ready to formulate our proposal for computing persistent homology of networks.

4. The Rips complex of a network

Recall that for a metric space (X, d_X), the *Rips complex* is defined for each $\delta \geqslant 0$ as follows:

$$\mathfrak{R}_X^\delta := \{\sigma \in \mathrm{pow}(X) : \mathrm{diam}(\sigma) \leqslant \delta\}, \text{ where } \mathrm{diam}(\sigma) := \max_{x,x' \in \sigma} d_X(x, x').$$

Following this definition, we can define the Rips complex for a network (X, ω_X) as follows:

$$\mathfrak{R}_X^\delta := \left\{\sigma \in \mathrm{pow}(X) : \max_{x,x' \in \sigma} \omega_X(x, x') \leqslant \delta\right\}.$$

We illustrate the Rips complex construction in Figure 2. To any network (X, ω_X), we may associate the *Rips filtration* $\{\mathfrak{R}_X^\delta \hookrightarrow \mathfrak{R}_X^{\delta'}\}_{\delta \leqslant \delta'}$. We denote the k-dimensional persistence vector space associated to this filtration by $\mathcal{H}_k^{\mathfrak{R}}(X)$, and the corresponding persistence diagram by $\mathrm{Dgm}_k^{\mathfrak{R}}(X)$. The Rips filtration is stable to small perturbations of the input data; this is the content of the next proposition.

Proposition 7. *Let* $(X, \omega_X), (Y, \omega_Y) \in \mathcal{N}$. *Then* $d_{\mathrm{B}}(\mathrm{Dgm}_k^{\mathfrak{R}}(X), \mathrm{Dgm}_k^{\mathfrak{R}}(Y)) \leqslant 2 d_{\mathcal{N}}(X, Y)$.

Proof. The crux of our proof lies in constructing diagrams of simplicial maps similar to those in Lemma 3 and checking that the appropriate pairs of maps are contiguous. Let $\eta = 2 d_{\mathcal{N}}(X, Y)$. Then by Proposition 4, there exist maps $\varphi : X \to Y$ and $\psi : Y \to X$ such that $\max(\mathrm{dis}(\varphi), \mathrm{dis}(\psi), C_{X,Y}(\varphi, \psi), C_{Y,X}(\psi, \varphi)) \leqslant \eta$. We begin by checking that φ, ψ induce simplicial maps $\varphi_\delta : \mathfrak{R}_X^\delta \to \mathfrak{R}_Y^{\delta+\eta}$ and $\psi_\delta : \mathfrak{R}_Y^\delta \to \mathfrak{R}_Y^{\delta+\eta}$ for each $\delta \in \mathbb{R}$.

Let $\delta' \geqslant \delta \in \mathbb{R}$. Let $\sigma = \{x_0, \ldots, x_n\} \in \mathfrak{R}_X^\delta$. Then $\omega_X(x_i, x_j) \leqslant \delta$ for each $0 \leqslant i, j \leqslant n$. Since $\mathrm{dis}(\varphi) \leqslant \eta$, we have the following for each i, j:

$$|\omega_X(x_i, x_j) - \omega_Y(\varphi(x_i), \varphi(x_j))| \leqslant \eta.$$

So $\omega_Y(\varphi(x_i), \varphi(x_j)) \leqslant \omega_X(x_i, x_j) + \eta \leqslant \delta + \eta$ for each $0 \leqslant i, j \leqslant n$. Thus $\varphi_\delta(\sigma) := \{\varphi(x_0), \ldots, \varphi(x_n)\}$ is a simplex in $\mathfrak{R}_Y^{\delta+\eta}$. Thus the map on simplices φ_δ induced by φ is simplicial for each $\delta \in \mathbb{R}$.

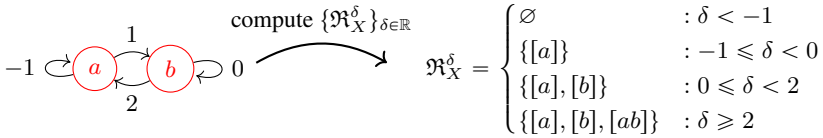
compute $\{\mathfrak{R}_X^\delta\}_{\delta \in \mathbb{R}}$
−1
1
a
b
0
2
$\mathfrak{R}_X^\delta = \begin{cases} \varnothing & : \delta < -1 \\ \{[a]\} & : -1 \leqslant \delta < 0 \\ \{[a], [b]\} & : 0 \leqslant \delta < 2 \\ \{[a], [b], [ab]\} & : \delta \geqslant 2 \end{cases}$

FIGURE 2. Computing the Rips complex of a network (X, ω_X).

Similarly we can check that the map ψ_δ on simplices induced by ψ is simplicial for each $\delta \in \mathbb{R}$. Now we have the following diagram of simplicial complexes and simplicial maps, for $\delta' \geqslant \delta \in \mathbb{R}$:

$$\begin{array}{ccccc} \mathfrak{R}_X^{\delta} & \xrightarrow{s_{\delta,\delta'}} & \mathfrak{R}_X^{\delta'} & & \\ & \searrow^{\varphi_\delta} & & \searrow^{\varphi_{\delta'}} & \\ & & \mathfrak{R}_Y^{\delta+\eta} & \xrightarrow{t_{\delta+\eta,\delta'+\eta}} & \mathfrak{R}_Y^{\delta'+\eta} \end{array}$$

Here $s_{\delta,\delta'}$ and $t_{\delta+\eta,\delta'+\eta}$ are the inclusion maps. We claim that $t_{\delta+\eta,\delta'+\eta} \circ \varphi_\delta$ and $\varphi_{\delta'} \circ s_{\delta,\delta'}$ are contiguous simplicial maps. To see this, let $\sigma \in \mathfrak{R}_X^{\delta}$. Since $s_{\delta,\delta'}$ is just the inclusion, it follows that $t_{\delta+\eta,\delta'+\eta}(\varphi_\delta(\sigma)) \cup \varphi_{\delta'}(s_{\delta,\delta'}(\sigma)) = \varphi_\delta(\sigma)$, which is a simplex in $\mathfrak{R}_Y^{\delta+\eta}$ because φ_δ is simplicial, and hence a simplex in $\mathfrak{R}_Y^{\delta'+\eta}$ because the inclusion $t_{\delta+\eta,\delta'+\eta}$ is simplicial. Thus $t_{\delta+\eta,\delta'+\eta} \circ \varphi_\delta$ and $\varphi_{\delta'} \circ s_{\delta,\delta'}$ are contiguous, and their induced linear maps upon passing to homology are equal. By a similar argument, and using the fact that $\mathrm{dis}(\psi) \leqslant \eta$, one can show that $s_{\delta+\eta,\delta'+\eta} \circ \psi_\delta$ and $\psi_{\delta'} \circ t_{\delta,\delta'}$ are contiguous simplicial maps as well, for each $\delta' \geqslant \delta \in \mathbb{R}$.

Next we check that the maps $\psi_{\delta+\eta} \circ \varphi_\delta$ and $s_{\delta,\delta+2\eta}$ in the figure below are contiguous.

$$\begin{array}{ccc} \mathfrak{R}_X^{\delta} & \xrightarrow{s_{\delta,\delta+2\eta}} & \mathfrak{R}_X^{\delta+2\eta} \\ & \searrow^{\varphi_\delta} \quad \nearrow_{\psi_{\delta+\eta}} & \\ & \mathfrak{R}_Y^{\delta+\eta} & \end{array}$$

Let $\sigma = [x_0, \ldots, x_n] \in \mathfrak{R}_X^{\delta}$. Then for any $x_i, x_j \in \sigma$, we have

$$\begin{aligned} |\omega_X(x_i, x_j) - \omega_X(\psi(\varphi(x_i)), \psi(\varphi(x_j)))| &\leqslant |\omega_X(x_i, x_j) - \omega_Y(\varphi(x_i), \varphi(x_j))| \\ &\quad + |\omega_Y(\varphi(x_i), \varphi(x_j)) - \omega_X(\psi(\varphi(x_i)), \psi(\varphi(x_j)))| \\ &\leqslant 2\eta. \end{aligned}$$

Thus we obtain $\omega_X(\psi(\varphi(x_i)), \psi(\varphi(x_j))) \leqslant \omega_X(x_i, x_j) + 2\eta \leqslant \delta + 2\eta$.

Since this holds for any $x_i, x_j \in \sigma$, it follows that $\psi_{\delta+\eta}(\varphi_\delta(\sigma)) \in \mathfrak{R}_X^{\delta+2\eta}$. We further claim that

$$\tau := \sigma \cup \psi_{\delta+\eta}(\varphi_\delta(\sigma)) = \{x_0, x_1, \ldots, x_n, \psi(\varphi(x_0)), \ldots, \psi(\varphi(x_n))\}$$

is a simplex in $\mathfrak{R}_X^{\delta+2\eta}$. Let $0 \leqslant i, j \leqslant n$. It suffices to show that $\omega_X(x_i, \psi(\varphi(x_j)) \leqslant \delta + 2\eta$.

Notice that because $C_{X,Y}(\varphi, \psi) \leqslant \eta$ (Proposition 4), we have

$$\max_{(x,y) \in X \times Y} |\omega_X(x, \psi(y)) - \omega_Y(\varphi(x), y)| \leqslant \eta.$$

Let $y = \varphi(x_j)$. Then $|\omega_X(x_i, \psi(y)) - \omega_Y(\varphi(x_i), y)| \leqslant \eta$. In particular,

$$\omega_X(x_i, \psi(\varphi(x_j))) \leqslant \omega_Y(\varphi(x_i), \varphi(x_j)) + \eta \leqslant \omega_X(x_i, x_j) + 2\eta \leqslant \delta + 2\eta.$$

Since $0 \leqslant i, j \leqslant n$ were arbitrary, it follows that $s_{\delta,\delta+2\eta}(\sigma) \cup \psi_{\delta+\eta}(\varphi_\delta(\sigma)) = \tau$ is a simplex in $\mathfrak{R}_X^{\delta+2\eta}$. Thus the maps $\psi_{\delta+\eta} \circ \varphi_\delta$ and $s_{\delta,\delta+2\eta}$ are contiguous. Similarly, one can show that $t_{\delta,\delta+2\eta}$ and $\varphi_{\delta+\eta} \circ \psi_\delta$ are contiguous. The result now follows by an application of Lemma 3. □

Remark 8. The preceding proposition serves a dual purpose: (1) it shows that the Rips persistence diagram is robust to noise in input data, and (2) it shows that instead of computing the network distance between two networks, one can compute the bottleneck distance between their Rips persistence diagrams as a suitable proxy. The advantage to computing bottleneck distance is that it can be done in polynomial time (see [EIK01]), whereas computing $d_{\mathcal{N}}$ is NP-hard in general [CM16b]. We remark that the idea of computing Rips persistence diagrams to compare finite metric spaces first appeared in [CCSG+09b], and moreover, that Proposition 7 is an extension of Theorem 3.1 in [CCSG+09b].

The Rips filtration in the setting of symmetric networks has been used in [HMR09, CH13, GPCI15, PSDV13], albeit without addressing stability results. To our knowledge, Proposition 7 is the first quantitative result justifying the constructions in these prior works.

5. The Dowker complex of a network

Given any network (X, ω_X), we can define an associated weight function $\overline{\omega}_X : X \times X \to \mathbb{R}$ as follows:

$$\overline{\omega}_X(x, x') := \max(\omega_X(x, x), \omega_X(x', x'), \omega_X(x, x')), \text{ for each } x, x' \in X.$$

Remark 9 ($\overline{\omega}_X$ is asymmetric)**.** It is important to note that $\overline{\omega}_X$ is still asymmetric, in general. Also, note that if ω_X happens to be a proper metric on X, then $\overline{\omega}_X = \omega_X$.

We have defined $\overline{\omega}_X$ for notational convenience, and use it in the next definition.

For any $\delta \in \mathbb{R}$, consider the following relation:

$$R_{\delta,X} := \left\{(x, x') : \overline{\omega}_X(x, x') \leqslant \delta\right\}. \tag{2}$$

Then $R_{\delta,X} \subseteq X \times X$, and $R_{\delta_F,X} = X \times X$ for some sufficiently large δ_F. Furthermore, for any $\delta' \geqslant \delta$, we have $R_{\delta,X} \subseteq R_{\delta',X}$. Using $R_{\delta,X}$, we build a simplicial complex $\mathfrak{D}^{\mathrm{si}}_{\delta}$ as follows:

$$\mathfrak{D}^{\mathrm{si}}_{\delta,X} := \left\{\sigma = [x_0, \ldots, x_n] : \text{ there exists } x' \in X \text{ such that } (x_i, x') \in R_{\delta,X} \text{ for each } x_i\right\}. \tag{3}$$

If $\sigma \in \mathfrak{D}^{\mathrm{si}}_{\delta,X}$, it is clear that any face of σ also belongs to $\mathfrak{D}^{\mathrm{si}}_{\delta,X}$. We call $\mathfrak{D}^{\mathrm{si}}_{\delta,X}$ the *Dowker δ-sink simplicial complex* associated to X, and refer to x' as a *δ-sink* for σ (where σ and x' should be clear from context).

Since $R_{\delta,X}$ is an increasing sequence of sets, it follows that $\mathfrak{D}^{\mathrm{si}}_{\delta,X}$ is an increasing sequence of simplicial complexes. In particular, for $\delta' \geqslant \delta$, there is a natural inclusion map $\mathfrak{D}^{\mathrm{si}}_{\delta,X} \hookrightarrow \mathfrak{D}^{\mathrm{si}}_{\delta',X}$. We write $\mathfrak{D}^{\mathrm{si}}_{X}$ to denote the filtration $\{\mathfrak{D}^{\mathrm{si}}_{\delta,X} \hookrightarrow \mathfrak{D}^{\mathrm{si}}_{\delta',X}\}_{\delta \leqslant \delta'}$ associated to X. We call this the *Dowker sink filtration on X*. We will denote the k-dimensional persistence diagram arising from this filtration by $\mathrm{Dgm}^{\mathrm{si}}_k(X)$.

Note that we can define a dual construction as follows:

$$\mathfrak{D}^{\mathrm{so}}_{\delta,X} := \left\{\sigma = [x_0, \ldots, x_n] : \text{ there exists } x' \in X \text{ such that } (x', x_i) \in R_{\delta,X} \text{ for each } x_i\right\}. \tag{4}$$

We call $\mathfrak{D}^{\mathrm{so}}_{\delta,X}$ the *Dowker δ-source simplicial complex* associated to X. The filtration $\{\mathfrak{D}^{\mathrm{so}}_{\delta,X} \hookrightarrow \mathfrak{D}^{\mathrm{so}}_{\delta',X}\}_{\delta \leqslant \delta'}$ associated to X is called the *Dowker source filtration*, denoted $\mathfrak{D}^{\mathrm{so}}_{X}$. We denote the k-dimensional persistence diagram arising from this filtration by $\mathrm{Dgm}^{\mathrm{so}}_k(X)$. Notice that any construction using $\mathfrak{D}^{\mathrm{si}}_{\delta,X}$ can also be repeated using $\mathfrak{D}^{\mathrm{so}}_{\delta,X}$, so we focus on the case of the sink complexes and restate results in terms of source complexes where necessary. In particular, we will prove in §5.1 that

$$\mathrm{Dgm}^{\mathrm{si}}_k(X) = \mathrm{Dgm}^{\mathrm{so}}_k(X) \text{ for any } k \in \mathbb{Z}_+,$$

so it makes sense to talk about "the" Dowker diagram associated to X.

As in the case of the Rips filtration, both the Dowker sink and source filtrations are stable. We state the next result in terms of sink filtrations, but a similar proof establishes an analogous result for source filtrations. Alternatively, the result for source filtrations will follow after we prove in §5.1 that both filtrations produce the same output persistence diagram.

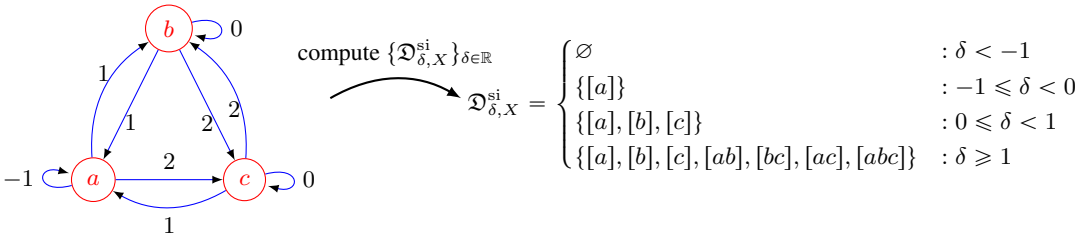

FIGURE 3. Computing the Dowker sink complex of a network (X, ω_X).

Proposition 10. *Let* $(X,\omega_X),(Y,\omega_Y)\in\mathcal{N}$. *Then* $d_{\mathrm{B}}(\mathrm{Dgm}_k^{\mathrm{si}}(X),\mathrm{Dgm}_k^{\mathrm{si}}(Y))\leqslant 2d_{\mathcal{N}}(X,Y)$.

Proof of Proposition 10. Let $\eta=2d_{\mathcal{N}}(X,Y)$. Then by Proposition 4, there exist maps $\varphi:X\to Y,\psi:Y\to X$ such that $\max(\mathrm{dis}(\varphi),\mathrm{dis}(\psi),C_{X,Y}(\varphi,\psi),C_{Y,X}(\psi,\varphi))\leqslant\eta$. First we check that φ,ψ induce simplicial maps $\varphi_\delta:\mathfrak{D}^{\mathrm{si}}_{\delta,X}\to\mathfrak{D}^{\mathrm{si}}_{\delta+\eta,Y}$ and $\psi_\delta:\mathfrak{D}^{\mathrm{si}}_{\delta,Y}\to\mathfrak{D}^{\mathrm{si}}_{\delta+\eta,Y}$ for each $\delta\in\mathbb{R}$.

Let $\delta'\geqslant\delta\in\mathbb{R}$. Let $\sigma=[x_0,\ldots,x_n]\in\mathfrak{D}^{\mathrm{si}}_{\delta,X}$. Then there exists $x'\in X$ such that $\overline{\omega}_X(x_i,x')\leqslant\delta$ for each $0\leqslant i\leqslant n$. Fix such an x'. Since $\mathrm{dis}(\varphi)\leqslant\eta$, we have the following for each i:

$$|\omega_X(x_i,x')-\omega_Y(\varphi(x_i),\varphi(x'))|\leqslant\eta.$$

Furthermore, for each i we have:

$$|\omega_X(x_i,x_i)-\omega_Y(\varphi(x_i),\varphi(x_i))|\leqslant\eta,$$
$$|\omega_X(x',x')-\omega_Y(\varphi(x'),\varphi(x'))|\leqslant\eta.$$

So $\overline{\omega}_Y(\varphi(x_i),\varphi(x'))\leqslant\overline{\omega}_X(x_i,x')+\eta\leqslant\delta+\eta$ for each $0\leqslant i\leqslant n$. Thus $\varphi_\delta(\sigma):=\{\varphi(x_0),\ldots,\varphi(x_n)\}$ is a simplex in $\mathfrak{D}^{\mathrm{si}}_{\delta+\eta,Y}$. Thus the map on simplices φ_δ induced by φ is simplicial for each $\delta\in\mathbb{R}$.

Similarly we can check that the map ψ_δ on simplices induced by ψ is simplicial. Now to prove the result, it will suffice to check the contiguity conditions in the statement of Lemma 3. Consider the following diagram:

$$\begin{array}{ccccc}
\mathfrak{D}^{\mathrm{si}}_{\delta,X} & \xrightarrow{s_{\delta,\delta'}} & \mathfrak{D}^{\mathrm{si}}_{\delta',X} & & \\
& \searrow^{\varphi_\delta} & & \searrow^{\varphi_{\delta'}} & \\
& & \mathfrak{D}^{\mathrm{si}}_{\delta+\eta,Y} & \xrightarrow{t_{\delta+\eta,\delta'+\eta}} & \mathfrak{D}^{\mathrm{si}}_{\delta'+\eta,Y}
\end{array}$$

Here $s_{\delta,\delta'}$ and $t_{\delta+\eta,\delta'+\eta}$ are the inclusion maps. We claim that $t_{\delta+\eta,\delta'+\eta}\circ\varphi_\delta$ and $\varphi_{\delta'}\circ s_{\delta,\delta'}$ are contiguous simplicial maps. To see this, let $\sigma\in\mathfrak{D}^{\mathrm{si}}_{\delta,X}$. Since $s_{\delta,\delta'}$ is just the inclusion, it follows that $t_{\delta+\eta,\delta'+\eta}(\varphi_\delta(\sigma))\cup\varphi_{\delta'}(s_{\delta,\delta'}(\sigma))=\varphi_\delta(\sigma)$, which is a simplex in $\mathfrak{D}^{\mathrm{si}}_{\delta+\eta,Y}$ because φ_δ is simplicial, and hence a simplex in $\mathfrak{D}^{\mathrm{si}}_{\delta'+\eta,Y}$ because the inclusion $t_{\delta+\eta,\delta'+\eta}$ is simplicial. Thus $t_{\delta+\eta,\delta'+\eta}\circ\varphi_\delta$ and $\varphi_{\delta'}\circ s_{\delta,\delta'}$ are contiguous, and their induced linear maps for homology are equal. By a similar argument, one can show that $s_{\delta+\eta,\delta'+\eta}\circ\psi_\delta$ and $\psi_{\delta'}\circ t_{\delta,\delta'}$ are contiguous simplicial maps as well.

Next we check that the maps $\psi_{\delta+\eta}\circ\varphi_\delta$ and $s_{\delta,\delta+2\eta}$ in the figure below are contiguous.

$$\begin{array}{ccc}
\mathfrak{D}^{\mathrm{si}}_{\delta,X} & \xrightarrow{s_{\delta,\delta+2\eta}} & \mathfrak{D}^{\mathrm{si}}_{\delta+2\eta,X} \\
& \searrow^{\varphi_\delta} \quad \nearrow_{\psi_{\delta+\eta}} & \\
& \mathfrak{D}^{\mathrm{si}}_{\delta+\eta,Y} &
\end{array}$$

Let $x_i\in\sigma$. Note that for our fixed $\sigma=[x_0,\ldots,x_n]\in\mathfrak{D}^{\mathrm{si}}_{\delta,X}$ and x', we have:

$$\begin{aligned}
|\omega_X(x_i,x')-\omega_X(\psi(\varphi(x_i)),\psi(\varphi(x')))| &\leqslant |\omega_X(x_i,x')-\omega_Y(\varphi(x_i),\varphi(x'))| \\
&\quad+|\omega_Y(\varphi(x_i),\varphi(x'))-\omega_X(\psi(\varphi(x_i)),\psi(\varphi(x')))| \\
&\leqslant 2\eta.
\end{aligned}$$

Thus we obtain $\omega_X(\psi(\varphi(x_i)),\psi(\varphi(x')))\leqslant\omega_X(x_i,x')+2\eta\leqslant\delta+2\eta$.

One can similarly obtain:

$$|\omega_X(x_i,x_i)-\omega_X(\psi(\varphi(x_i)),\psi(\varphi(x_i)))|\leqslant 2\eta,$$
$$|\omega_X(x',x')-\omega_X(\psi(\varphi(x')),\psi(\varphi(x')))|\leqslant 2\eta.$$

It then follows that $\overline{\omega}_X(\psi(\varphi(x_i)),\psi(\varphi(x'))) \leqslant \overline{\omega}_X(x_i,x') + 2\eta \leqslant \delta + 2\eta$. Since this holds for any $x_i \in \sigma$, it follows that $\psi_{\delta+\eta}(\varphi_\delta(\sigma)) \in \mathfrak{D}^{\mathrm{si}}_{\delta+2\eta,X}$. We further claim that

$$\tau := \sigma \cup \psi_{\delta+\eta}(\varphi_\delta(\sigma)) = \{x_0, x_1, \ldots, x_n, \psi(\varphi(x_0)), \ldots, \psi(\varphi(x_n))\}$$

is a simplex in $\mathfrak{R}_X^{\delta+2\eta}$. Let $0 \leqslant i \leqslant n$. It suffices to show that $\overline{\omega}_X(x_i, \psi(\varphi(x')) \leqslant \delta + 2\eta$.

Notice that from the reformulation of $d_{\mathcal{N}}$ (Proposition 4), we have

$$\max_{(x,y)\in X\times Y} |\omega_X(x,\psi(y)) - \omega_Y(\varphi(x),y)| \leqslant \eta.$$

Let $y = \varphi(x')$. Then $|\omega_X(x_i,\psi(y)) - \omega_Y(\varphi(x_i),y)| \leqslant \eta$. In particular,

$$\omega_X(x_i,\psi(\varphi(x'))) \leqslant \omega_Y(\varphi(x_i),\varphi(x')) + \eta \leqslant \omega_X(x_i,x') + 2\eta \leqslant \delta + 2\eta.$$

Also note that $\omega_X(x_i,x_i) \leqslant \delta$, and $\omega_X(\psi(\varphi(x')),\psi(\varphi(x'))) \leqslant \delta + 2\eta$, by what we have already shown. Thus $\overline{\omega}_X(x_i,\psi(\varphi(x'))) \leqslant \delta + 2\eta$.

Since $0 \leqslant i \leqslant n$ were arbitrary, it follows that $\tau \in \mathfrak{D}^{\mathrm{si}}_{\delta+2\eta,X}$. Thus the maps $\psi_{\delta+\eta}\circ\varphi_\delta$ and $s_{\delta,\delta+2\eta}$ are contiguous. Similarly, one can show that $t_{\delta,\delta+2\eta}$ and $\varphi_{\delta+\eta}\circ\psi_\delta$ are contiguous.

The result now follows by an application of Lemma 3. □

Remark 11. The preceding proposition shows that the Dowker persistence diagram is robust to noise in input data, and that the bottleneck distance between Dowker persistence diagrams arising from two networks can be used as a proxy for computing the actual network distance. Note the analogy with Remark 8.

Both the Dowker and Rips filtrations are valid methods for computing persistent homology of networks, by virtue of their stability results (Propositions 7 and 10). However, we present the Dowker filtration as an appropriate method for capturing directionality information in directed networks. In §5.2 we discuss this particular feature of the Dowker filtration in more detail.

Remark 12 (Dissimilarity networks)**.** A *dissimilarity network* is a network (X, A_X) where $A_X : X \times X \to [0,\infty)$ is a *dissimilarity function*, i.e. a map such that $A_X(x,x') = 0 \iff x = x'$, for any $x, x' \in X$. Thus (X, A_X) deserves to be treated as a weighted, directed network whose weight matrix vanishes precisely on the diagonal. The collection of all such networks will be denoted $\mathcal{N}^{\mathrm{dis}}$.

In the case of dissimilarity networks, the definition of a Dowker complex becomes simpler. For each $\delta \geqslant 0$, let R_δ denote the relation $\{(x,x') : A_X(x,x') \leqslant \delta\}$. Then $R_\delta \subseteq X \times X$, $R_0 = \{(x,x) : x \in X\}$ and $R_{\delta_F} = X \times X$ for some sufficiently large δ_F. The Dowker sink and source complexes are then defined as before.

Remark 13 (Symmetric networks)**.** In the simplified setting of symmetric networks, the Dowker sink and source filtrations coincide, and so we automatically obtain $\mathrm{Dgm}_k^{\mathrm{so}}(X) = \mathrm{Dgm}_k^{\mathrm{si}}(X)$ for any $k \in \mathbb{Z}_+$ and any $(X,\omega_X) \in \mathcal{N}$.

Remark 14 (The metric space setting and relation to witness complexes)**.** When restricted to the setting of metric spaces, the Dowker complex resembles a construction called the witness complex [DSC04]. In particular, a version of the Dowker complex for metric spaces, constructed in terms of *landmarks* and *witnesses*, was discussed in [CDSO14], along with stability results. When restricted to the special networks that are metric spaces, our definitions and results agree with those presented in [CDSO14].

5.1. **Dowker duality and equivalence of diagrams.** Let $(X,\omega_X) \in \mathcal{N}$. In the preceding section, we have provided the constructions of the vector spaces $H_k(\mathfrak{D}^{\mathrm{si}}_{\delta,X})$ and $H_k(\mathfrak{D}^{\mathrm{so}}_{\delta,X})$, for any $k \in \mathbb{Z}_+$. By a theorem of Dowker [Dow52], these vector spaces are actually isomorphic:

Theorem 15 (Dowker (1952))**.** *Let $(X,\omega_X) \in \mathcal{N}$, let $\delta \in \mathbb{R}$, and let $k \in \mathbb{Z}_+$. Then,*

$$H_k(\mathfrak{D}^{\mathrm{si}}_{\delta,X}) \cong H_k(\mathfrak{D}^{\mathrm{so}}_{\delta,X}).$$

In the modern language of persistent homology, the more interesting result would be to show that for any dimension $k \in \mathbb{Z}_+$, the persistence diagrams of the Dowker sink and source filtrations are equal. Such a result appears to be known in the applied algebraic topology community (see [CDSO14] for a mention of this result, which we call *Dowker duality*), but we were unable to find a proof in any published work. In this section, we provide a detailed proof of this result.

As a first step, we state the Persistence Equivalence Theorem [EH10].

Theorem 16 (Persistence Equivalence Theorem)**.** *Consider two persistence vector spaces* $\mathcal{U} = \{U^{\delta_i} \xrightarrow{\mu_{i,i+1}} U^{\delta_{i+1}}\}_{i\in\mathbb{N}}$ *and* $\mathcal{V} = \{V^{\delta_i} \xrightarrow{\nu_{i,i+1}} V^{\delta_{i+1}}\}_{i\in\mathbb{N}}$ *with connecting maps* $\varphi_i : U^{\delta_i} \to V^{\delta_i}$.

$$\begin{array}{ccccccccc}
\cdots & \longrightarrow & V^{\delta_i} & \longrightarrow & V^{\delta_{i+1}} & \longrightarrow & V^{\delta_{i+2}} & \longrightarrow & \cdots \\
 & & \uparrow{\scriptstyle \varphi_i} & & \uparrow{\scriptstyle \varphi_{i+1}} & & \uparrow{\scriptstyle \varphi_{i+2}} & & \\
\cdots & \longrightarrow & U^{\delta_i} & \longrightarrow & U^{\delta_{i+1}} & \longrightarrow & U^{\delta_{i+2}} & \longrightarrow & \cdots
\end{array}$$

If the φ_i *are all isomorphisms and each square in the diagram above commutes, then:*

$$\mathrm{Dgm}(\mathcal{U}) = \mathrm{Dgm}(\mathcal{V}).$$

Theorem 17 (Dowker duality)**.** *Let* $(X, \omega_X) \in \mathcal{N}$*, and let* $k \in \mathbb{Z}_+$*. Then,*

$$\mathrm{Dgm}_k^{\mathrm{si}}(X) = \mathrm{Dgm}_k^{\mathrm{so}}(X).$$

Thus we may call either of the two diagrams above the k-dimensional Dowker diagram of X*, denoted* $\mathrm{Dgm}_k^{\mathfrak{D}}(X)$.

Before proving the theorem, we recall the construction of a combinatorial barycentric subdivision, and also a special chain map called the *inverse chain derivation*. These constructions are attributed to Lefschetz [Lef42, §4.7], and a detailed treatment using standard notation appears in [Cro78, §7].

Definition 1. For any simplicial complex Σ, one may construct a new simplicial complex $\Sigma^{(1)}$, called the *first barycentric subdivision*, as follows:

$$\Sigma^{(1)} := \{[\sigma_1, \sigma_2, \ldots, \sigma_p] : \sigma_1 \subseteq \sigma_2 \subseteq \ldots \subseteq \sigma_p, \text{ each } \sigma_i \in \Sigma\}.$$

Note that the vertices of $\Sigma^{(1)}$ are the simplices of Σ, and the simplices of $\Sigma^{(1)}$ are nested sequences of simplices of Σ. Furthermore, note that given any two simplicial complexes Σ, Ξ and a simplicial map $f : \Sigma \to \Xi$, there is a natural simplicial map $f^{(1)} : \Sigma^{(1)} \to \Xi^{(1)}$ defined as:

$$f^{(1)}([\sigma_1, \ldots, \sigma_p]) := [f(\sigma_1), \ldots, f(\sigma_p)], \qquad \sigma_1 \subseteq \sigma_2 \subseteq \ldots, \sigma_p, \text{ each } \sigma_i \in \Sigma.$$

To see that this is simplicial, note that $f(\sigma_i) \subseteq f(\sigma_j)$ whenever $\sigma_i \subseteq \sigma_j$. Finally, observe that the inclusion map $\iota : \Sigma \to \Sigma$ induces an inclusion map $\iota^{(1)} : \Sigma^{(1)} \to \Sigma^{(1)}$.

Let Σ be a simplicial complex. Since all the simplicial complexes we deal with are finite, we may assume without loss of generality that the vertices of Σ are totally ordered, so that the vertices of any simplex are totally ordered. One may then define an associated map $\varphi_\Sigma : \Sigma^{(1)} \to \Sigma$ as follows: first define φ_Σ on vertices of $\Sigma^{(1)}$ by $\varphi_\Sigma(\sigma) = s_\sigma$, where s_σ is the least vertex of σ with respect to the total order. Next, for any simplex $[\sigma_1, \ldots, \sigma_p]$ of $\Sigma^{(1)}$, where $\sigma_1 \subseteq \ldots \subseteq \sigma_p$, we have $\varphi_\Sigma(\sigma_i) = s_{\sigma_i} \in \sigma_p$ for all $1 \leqslant i \leqslant p$. Thus $[\varphi_\Sigma(\sigma_1), \ldots, \varphi_\Sigma(\sigma_p)] = [s_{\sigma_1}, s_{\sigma_2}, \ldots, s_{\sigma_p}]$ is a face of σ_p, hence a simplex of Σ. This defines φ_Σ as a simplicial map $\Sigma^{(1)} \to \Sigma$.

Since φ_Σ is simplicial, it induces a chain map $(\varphi_\Sigma)_*$ between the corresponding chain complexes [Mun84, §1.12]. This chain map is called the *inverse chain derivation*. It satisfies the special property that the induced map $(\varphi_\Sigma)_\# : H_\bullet(\Sigma^{(1)}) \to H_\bullet(\Sigma)$ between the corresponding homology vector spaces is an isomorphism [Cro78, Theorem 7.5].

Proof of Theorem 17. Let $(X, \omega_X) \in \mathcal{N}$, let $k \in \mathbb{Z}_+$, and let $(\delta_i)_{i=1}^n$ be an increasing sequence of real numbers. For each δ_i, we have the corresponding relation $R_{\delta_i,X}$ given by Equation 2. For notational convenience, we will henceforth write $R_{\delta_i} := R_{\delta_i,X}$. Notice that we then have the following filtered complexes:

$$\varnothing \subseteq \mathfrak{D}^{\text{so}}_{\delta_1,X} \subseteq \mathfrak{D}^{\text{so}}_{\delta_2,X} \subseteq \dots \mathfrak{D}^{\text{so}}_{\delta_n,X}$$

$$\varnothing \subseteq \mathfrak{D}^{\text{si}}_{\delta_1,X} \subseteq \mathfrak{D}^{\text{si}}_{\delta_2,X} \subseteq \dots \subseteq \mathfrak{D}^{\text{si}}_{\delta_n,X}.$$

By Dowker's theorem (15), $H_k(\mathfrak{D}^{\text{si}}_{\delta_i,X}) \cong H_k(\mathfrak{D}^{\text{so}}_{\delta_i,X})$ for each $i \in \mathbb{N}$, via an isomorphism ω_i. We need to show that the Dowker sink and source diagrams are equal. To see this, let $n_1, n_2 \in \mathbb{N}$, $n_1 < n_2$, and write:

$$E_1 := \mathfrak{D}^{\text{si}}_{\delta_{n_1},X},\ E_2 := \mathfrak{D}^{\text{si}}_{\delta_{n_2},X},\ F_1 := \mathfrak{D}^{\text{so}}_{\delta_{n_1},X},\ F_2 := \mathfrak{D}^{\text{so}}_{\delta_{n_2},X}.$$

Note that the simplicial inclusion maps $\iota_E : E_1 \to E_2$ and $\iota_F : F_1 \to F_2$ induce linear maps $(\iota_E)_\# : H_k(E_1) \to H_k(E_2)$ and $(\iota_F)_\# : H_k(F_1) \to H_k(F_2)$ at the homology level. We will proceed by constructing isomorphisms $\omega_1 : H_k(E_1) \to H_k(F_1)$ and $\omega_2 : H_k(E_2) \to H_k(F_2)$, and then showing that the following diagram commutes:

$$\begin{array}{ccc} H_k(E_1) & \xrightarrow{(\iota_E)_\#} & H_k(E_2) \\ \downarrow{\scriptstyle \omega_1} & & \downarrow{\scriptstyle \omega_2} \\ H_k(F_1) & \xrightarrow{(\iota_F)_\#} & H_k(F_2) \end{array}$$

Since $n_2 > n_1$ were arbitrary, we can then apply Theorem 16 to obtain the equivalence of diagrams that we need.

Let $\sigma = [x_1, \dots, x_p]$ be a simplex of E_1. Then there exists $x'_\sigma \in X$ such that $(x_i, x'_\sigma) \in R_{\delta_{n_1}}$ for all $1 \leqslant i \leqslant p$. For each $\sigma \in E_1$, we fix a choice of $x'_\sigma \in X$ such that $(x, x'_\sigma) \in R_{\delta_{n_1}}$ for each $x \in \sigma$.

Define $\psi_{F_1} : E_1^{(1)} \to F_1$ as follows: first, for any vertex $\sigma \in E_1^{(1)}$, define $\psi_{F_1}(\sigma) := x'_\sigma$. Note that $\psi_{F_1}(\sigma)$ is then a sink for σ. Next, for any simplex $[\sigma_1, \sigma_2, \dots, \sigma_q] \in E_1^{(1)}$, define

$$\psi_{F_1}([\sigma_1, \sigma_2, \dots, \sigma_q]) := [x'_{\sigma_1}, x'_{\sigma_2}, \dots, x'_{\sigma_q}].$$

Let $x \in \sigma_1$. Since $\sigma_1 \subseteq \sigma_2 \subseteq \dots \subseteq \sigma_q$, we have $x \in \sigma_i$ for each $1 \leqslant i \leqslant q$. Thus $(x, x'_{\sigma_1}), \dots, (x, x'_{\sigma_q}) \in R_{\delta_{n_1}}$ for each $1 \leqslant i \leqslant q$. Thus $[x'_{\sigma_1}, x'_{\sigma_2}, \dots, x'_{\sigma_q}] \in F_1$, with x as the source. This shows that ψ_{F_1} is a simplicial map. Similarly define $\psi_{F_2} : E_2^{(1)} \to F_2$. From the discussion preceding this proof, we also have simplicial maps $\varphi_{E_1} : E_1^{(1)} \to E_1$ and $\varphi_{E_2} : E_2^{(1)} \to E_2$ that induce inverse chain derivations. We also have simplicial inclusion maps $\iota_E : E_1 \to E_1, \iota_F : F_1 \to F_2, \iota_{E^{(1)}} : E_1^{(1)} \to E_2^{(1)}$. Now consider the following diagram of simplicial complexes and simplicial maps:

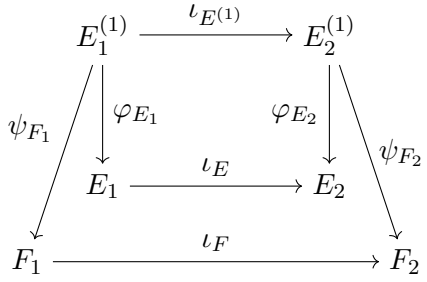

Fact. Consider the following linear maps:

$$\omega_1 := (\psi_{F_1})_\# \circ (\varphi_{E_1})_\#^{-1}, \qquad \omega_2 := (\psi_{F_2})_\# \circ (\varphi_{E_2})_\#^{-1}.$$

In Dowker's original proof of Theorem 15, $\omega_1 : H_k(E_1) \to H_k(F_1)$ and $\omega_2 : H_k(E_2) \to H_k(F_2)$ are shown to be isomorphisms [Dow52, Theorem 1a], a fact we use without reproving.

Claim 1. $\iota_E \circ \varphi_{E_1}$ and $\varphi_{E_2} \circ \iota_{E^{(1)}}$ are contiguous.

Claim 2. $\iota_F \circ \psi_{F_1}$ and $\psi_{F_2} \circ \iota_{E^{(1)}}$ are contiguous.

Suppose both claims are true. Recall that contiguous simplicial maps induce equal linear maps upon passing to homology. Thus we obtain:

$$\begin{aligned}
(\iota_F)_\# \circ (\psi_{F_1})_\# &= (\iota_F \circ \psi_{F_1})_\# && \text{(by functoriality of homology)}\\
&= (\psi_{F_2} \circ \iota_{E^{(1)}})_\# && \text{(by Claim 2)}\\
&= (\psi_{F_2})_\# \circ (\iota_{E^{(1)}})_\# && \text{(by functoriality of homology)}\\
&= (\psi_{F_2})_\# \circ (\varphi_{E_2})_\#^{-1} \circ (\iota_E)_\# \circ (\varphi_{E_1})_\# && \text{(by Claim 1 and functoriality)}\\
(\iota_F)_\# \circ (\psi_{F_1})_\# \circ (\varphi_{E_1})_\#^{-1} &= (\psi_{F_2})_\# \circ (\varphi_{E_2})_\#^{-1} \circ (\iota_E)_\# && \text{(by taking an inverse on the right).}
\end{aligned}$$

But now we have $(\iota_F)_\# \circ \omega_1 = \omega_2 \circ (\iota_E)_\#$, which is the commutativity relation we needed. The situation is summarized in the following diagram of vector spaces and linear maps:

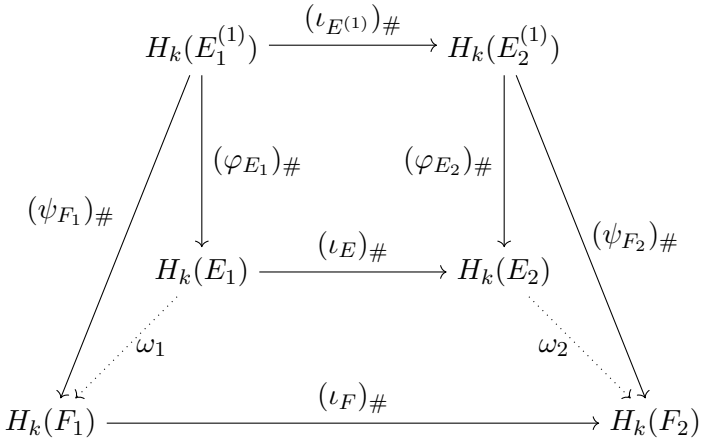

It remains to show that Claims 1 and 2 are true.

Proof of Claim 1. Let $\sigma' = [\sigma_0, \ldots, \sigma_p]$ be a simplex in $E_1^{(1)}$. Then σ_i is a simplex in E_1 for each $0 \leqslant i \leqslant p$, and $\sigma_0 \subseteq \ldots \subseteq \sigma_p$. Note that $\varphi_{E_1}(\sigma_i) \in \sigma_i \subseteq \sigma_p$ for each $0 \leqslant i \leqslant p$. Thus $\iota_E(\varphi_{E_1}(\sigma_i)) \in \sigma_p$ for each $0 \leqslant i \leqslant p$. Similarly $\varphi_{E_2}(\iota_{E^{(1)}}(\sigma_i)) = \varphi_{E_2}(\sigma_i) \in \sigma_p$ for each $0 \leqslant i \leqslant p$. Thus we have:

$$\begin{aligned}
\varphi_{E_2}(\iota_{E^{(1)}}(\sigma')) &= [\varphi_{E_2}(\sigma_1), \varphi_{E_2}(\sigma_2), \ldots, \varphi_{E_2}(\sigma_p)] \subseteq \sigma_p\\
\iota_E(\varphi_{E_1}(\sigma')) &= [\varphi_{E_1}(\sigma_1), \varphi_{E_1}(\sigma_2), \ldots, \varphi_{E_1}(\sigma_p)] \subseteq \sigma_p.
\end{aligned}$$

Thus $\varphi_{E_2}(\iota_{E^{(1)}}(\sigma'))$ and $\iota_E(\varphi_{E_1}(\sigma'))$ are faces of the same simplex. This holds for arbitrary $\sigma' \in E_1^{(1)}$. Thus the two maps are contiguous. This proves Claim 1. ∎

Proof of Claim 2. Let $\sigma' = [\sigma_0, \ldots, \sigma_p]$ be a simplex in $E_1^{(1)}$. Note that $\sigma_0 \subseteq \ldots \subseteq \sigma_p$. Let $x_0 \in \sigma_0$. Then $x_0 \in \sigma_i$ for each $0 \leqslant i \leqslant p$. Recall that $\psi_{F_2}(\sigma_i)$ is a sink for σ_i. Thus for each $0 \leqslant i \leqslant p$, we have

$$\big(x_0, (\psi_{F_2} \circ \iota_{E^{(1)}})(\sigma_i)\big) = \big(x_0, (\psi_{F_2}) \circ (\iota_{E^{(1)}})(\sigma_i)\big) = \big(x_0, \psi_{F_2}(\iota_{E^{(1)}}(\sigma_i))\big) = \big(x_0, \psi_{F_2}(\sigma_i)\big) \in R_{\delta_n},$$

where the first equality holds by functoriality of homology, and the last equality follows because inclusion maps induce inclusion maps upon taking barycentric subdivisions (Definition 1).

Similarly, we have $(x_0, \psi_{F_1}(\sigma_i)) \in R_{\delta_{n_1}}$ for each $0 \leqslant i \leqslant p$. Thus we obtain, for any $0 \leqslant i \leqslant p$,

$$\big(x_0, (\iota_F \circ \psi_{F_1})(\sigma_i)\big) = \big(x_0, \psi_{F_1}(\sigma_i)\big) \in R_{\delta_{n_1}} \subseteq R_{\delta_{n_2}},$$

where the first equality follows from functoriality of homology as before, and the last inclusion holds because $n_1 < n_2$, i.e. $\delta_{n_1} \leqslant \delta_{n_2}$ (Equation 2). Thus $\iota_F(\psi_{F_1}(\sigma')) \cup \psi_{F_2}(\iota_{E^{(1)}}(\sigma'))$ is a simplex in F_2 with x_0 as a source. Since $\sigma' \in E_1^{(1)}$ was arbitrary, the two maps are contiguous. This proves Claim 2. ■

Our result now follows by an application of the Persistence Equivalence Theorem. □

5.2. **Dowker persistence diagrams capture asymmetry.** From the very definitions of the Dowker source or sink complexes and the Rips complex at any given resolution, one can see that the Rips complex is blind to asymmetry in the input data, whereas either of the Dowker complexes is sensitive to asymmetry. Thus when analyzing datasets containing asymmetric information, one may wish to use the Dowker filtration instead of the Rips filtration. In particular, this property suggests that the Dowker persistence diagram is a stronger invariant for directed networks than the Rips persistence diagram. In this section, we provide a family of examples, called *cycle networks*, for which the Dowker persistence diagrams capture meaningful structure, whereas the Rips persistence diagrams do not. As a warm-up to analyzing this particular family of examples, we first probe the question "What happens to the Dowker or Rips persistence diagram of a network upon reversal of one (or more) edges?" Intuitively, if either of these persistence diagrams captures asymmetry, we would see a change in the diagram after applying this reversal operation to an edge. More concretely, we can make the following definition.

Definition 2 (Pair swaps)**.** Let $(X, \omega_X) \in \mathcal{N}$ be a network. For any $z, z' \in X$, define the (z, z')-*swap* of (X, ω_X) to be the network $S_X(z, z') := (X^{z,z'}, \omega_X^{z,z'})$ defined as follows:

$$X^{z,z'} := X,$$

$$\text{For any } x, x' \in X^{z,z'}, \qquad \omega_X^{z,z'}(x, x') := \begin{cases} \omega_X(x', x) & : x = z, x' = z' \\ \omega_X(x', x) & : x' = z, x = z' \\ \omega_X(x, x') & : \text{otherwise.} \end{cases}$$

In this language, the question of interest can be stated as follows: Given a network (X, ω_X) and a (x, x')-swap $S_X(x, x')$ for any $x, x' \in X$, how do the Rips or Dowker persistence diagrams of $S_X(x, x')$ differ from those of (X, ω_X)? This situation is illustrated in Figure 4. Example 21 shows an example where the Dowker persistence diagram captures the variation in a network that occurs after a pair swap, whereas the Rips persistence diagram fails to capture this difference. Furthermore, Remark 19 shows that Rips persistence diagrams always fail to do so.

We also consider the extreme situation where all the directions of the edges of a network are reversed, i.e. the network obtained by applying the pair swap operation to each pair of nodes. We would intuitively expect that the persistence diagrams would not change. The following discussion shows that the Rips and Dowker persistence diagrams are invariant under taking the transpose of a network.

Proposition 18. *Let $(X, \omega_X) \in \mathcal{N}$, and let $X^\top$ denote its transpose, i.e. the network $(X, \omega_X^\top)$ where $\omega_X^\top(x, x') := \omega_X(x', x)$ for $x, x' \in X$. Let $k \in \mathbb{Z}_+$. Then $\mathrm{Dgm}_k^{\mathrm{si}}(X) = \mathrm{Dgm}_k^{\mathrm{so}}(X^\top)$, and therefore $\mathrm{Dgm}_k^{\mathfrak{D}}(X) = \mathrm{Dgm}_k^{\mathfrak{D}}(X^\top)$. Moreover, we have $\mathrm{Dgm}_k^{\mathfrak{R}}(X) = \mathrm{Dgm}_k^{\mathfrak{R}}(X^\top)$.*

Proof of Proposition 18. Let $\delta \in \mathbb{R}$. We first claim that $\mathfrak{D}_\delta^{\mathrm{si}}(X) = \mathfrak{D}_\delta^{\mathrm{so}}(X^\top)$. Let $\sigma \in \mathfrak{D}_\delta^{\mathrm{si}}(X)$. Then there exists x' such that $\overline{\omega}_X(x, x') \leqslant \delta$ for any $x \in \sigma$. Thus $\overline{\omega}_{X^\top}(x', x) \leqslant \delta$. So $\sigma \in \mathfrak{D}_\delta^{\mathrm{so}}(X^\top)$. A similar argument shows the reverse containment. This proves our claim. Thus for $\delta \leqslant \delta' \leqslant \delta''$, we obtain the

following diagram:

$$\begin{array}{ccccccc} \mathfrak{D}^{\mathrm{si}}_{\delta}(X) & \longrightarrow & \mathfrak{D}^{\mathrm{si}}_{\delta'}(X) & \longrightarrow & \mathfrak{D}^{\mathrm{si}}_{\delta''}(X) & \longrightarrow & \dots \\ \| & & \| & & \| & & \\ \mathfrak{D}^{\mathrm{so}}_{\delta}(X^{\top}) & \longrightarrow & \mathfrak{D}^{\mathrm{so}}_{\delta'}(X^{\top}) & \longrightarrow & \mathfrak{D}^{\mathrm{so}}_{\delta''}(X^{\top}) & \longrightarrow & \dots \end{array}$$

Since the maps $\mathfrak{D}^{\mathrm{si}}_{\delta} \to \mathfrak{D}^{\mathrm{si}}_{\delta'}$, $\mathfrak{D}^{\mathrm{so}}_{\delta} \to \mathfrak{D}^{\mathrm{so}}_{\delta'}$ for $\delta' \geqslant \delta$ are all inclusion maps, it follows that the diagrams commute. Thus at the homology level, we obtain, via functoriality of homology, a commutative diagram of vector spaces where the intervening vertical maps are isomorphisms. By the Persistence Equivalence Theorem (16), the diagrams $\mathrm{Dgm}^{\mathrm{si}}_k(X)$ and $\mathrm{Dgm}^{\mathrm{so}}_k(X^{\top})$ are equal. By invoking Theorem 17, we obtain $\mathrm{Dgm}^{\mathfrak{D}}_k(X) = \mathrm{Dgm}^{\mathfrak{D}}_k(X^{\top})$.

Next note that for any $\sigma \in \mathrm{pow}(X)$, we have:

$$\max_{x,x' \in \sigma} \omega_X(x,x') = \max_{x,x' \in \sigma} \omega_X^{\top}(x,x').$$

From this observation, one can apply the Persistence Equivalence Theorem as above to see that the Rips persistence diagram of a network and its transpose are the same. □

Remark 19 (Pair swaps and their effect)**.** Let $(X,\omega_X) \in \mathcal{N}$, let $z,z' \in X$, and let $\sigma \in \mathrm{pow}(X)$. Then we have:

$$\max_{x,x' \in \sigma} \omega_X(x,x') = \max_{x,x' \in \sigma} \omega_X^{z,z'}(x,x').$$

Using this observation, one can then repeat the arguments used in the proof of Proposition 18 to show that:

$$\mathrm{Dgm}^{\mathfrak{R}}_k(X) = \mathrm{Dgm}^{\mathfrak{R}}_k(S_X(z,z')), \text{ for each } k \in \mathbb{Z}_+.$$

This encodes the intuitive fact that Rips persistence diagrams are blind to pair swaps.

On the other hand, k-dimensional Dowker persistence diagrams are not necessarily invariant to pair swaps when $k \geqslant 1$. Indeed, Example 21 constructs a space X for which there exist points $z,z' \in X$ such that

$$\mathrm{Dgm}^{\mathfrak{D}}_1(X) \neq \mathrm{Dgm}^{\mathfrak{D}}_1(S_X(z,z')).$$

However, 0-dimensional Dowker persistence diagrams are still invariant to pair swaps, as we show below.

Proposition 20. *Let $(X,\omega_X) \in \mathcal{N}$, let $z,z' \in X$, and let $\sigma \in \mathrm{pow}(X)$. Then we have:*

$$\mathrm{Dgm}^{\mathfrak{D}}_0(X) = \mathrm{Dgm}^{\mathfrak{D}}_0(S_X(z,z')).$$

Proof of Proposition 20. Let $\delta \in \mathbb{R}$. For notational convenience, we write, for each $k \in \mathbb{Z}_+$,

$$\begin{array}{lll} \mathfrak{D}^{\mathrm{si}}_{\delta} := \mathfrak{D}^{\mathrm{si}}_{\delta,X} & C^{\delta}_k := C_k(\mathfrak{D}^{\mathrm{si}}_{\delta,X}) & \partial^{\delta}_k := \partial^{\delta}_k : C^{\delta}_k \to C^{\delta}_{k-1} \\ \mathfrak{D}^{\mathrm{si}}_{\delta,S} := \mathfrak{D}^{\mathrm{si}}_{\delta,S_X(z,z')} & C^{\delta,S}_k := C_k(\mathfrak{D}^{\mathrm{si}}_{\delta,S_X(z,z')}) & \partial^{\delta,S}_k := \partial^{\delta,S}_k : C^{\delta,S}_k \to C^{\delta,S}_{k-1}. \end{array}$$

First note that pair swaps do not affect the entry of 0-simplices into the Dowker filtration. More precisely, for any $x \in X$, we can unpack the definition of $R_{\delta,X}$ (Equation 2) to obtain:

$$[x] \in \mathfrak{D}^{\mathrm{si}}_{\delta} \iff \omega_X(x,x) \leqslant \delta \iff \omega_X^{z,z'}(x,x) \leqslant \delta \iff [x] \in \mathfrak{D}^{\mathrm{si}}_{\delta,S}.$$

Thus for any $\delta \in \mathbb{R}$, we have $C^{\delta}_0 = C^{\delta,S}_0$. Since all 0-chains are automatically 0-cycles, we have $\ker(\partial^{\delta}_0) = \ker(\partial^{\delta,S}_0)$.

Next we wish to show that $\mathrm{im}(\partial^{\delta}_1) = \mathrm{im}(\partial^{\delta,S}_1)$ for each $\delta \in \mathbb{R}$. Let $\gamma \in C^{\delta}_1$. We first need to show the forward inclusion, i.e. that $\partial^{\delta}_1(\gamma) \in \mathrm{im}(\partial^{\delta,S}_1)$. It suffices to show this for the case that γ is a single 1-simplex $[x,x'] \in \mathfrak{D}^{\mathrm{si}}_{\delta}$; the case where γ is a linear combination of 1-simplices will then follow by linearity. Let $\gamma = [x,x'] \in \mathfrak{D}^{\mathrm{si}}_{\delta}$ for $x,x' \in X$. Then we have the following possibilities:

(1) $x'' \in X \backslash \{z,z'\}$ is a δ-sink for $[x,x']$.
(2) z (or z') is the only δ-sink for $[x,x']$, and $x,x' \notin \{z,z'\}$.

(3) z (or z') is the only δ-sink for $[x, x']$, and either x or x' belongs to $\{z, z'\}$.
(4) z (or z') is the only δ-sink for $[x, x']$, and both x, x' belong to $\{z, z'\}$.

In cases (1), (2), and (4), the (z, z')-pair swap has no effect on $[x, x']$, in the sense that we still have $[x, x'] \in \mathfrak{D}^{\mathrm{si}}_{\delta,S}$. So $[x'] - [x] = \partial^{\delta}_1(\gamma) = \partial^{\delta,S}_1(\gamma) \in \mathrm{im}(\partial^{\delta,S}_1)$. Next consider case (3), and assume for notational convenience that $[x, x'] = [z, x']$ and z' is the only δ-sink for $[z, x']$. By the definition of a δ-sink, we have $\overline{\omega}_X(z, z') \leqslant \delta$ and $\overline{\omega}_X(x', z') \leqslant \delta$. Notice that we also have:

$$[z, z'], [z', x'] \in \mathfrak{D}^{\mathrm{si}}_{\delta}, \text{ with } z' \text{ as a } \delta\text{-sink.}$$

After the (z, z')-pair swap, we still have $\overline{\omega}^{z,z'}_X(x', z') \leqslant \delta$, but possibly $\overline{\omega}^{z,z'}_X(z, z') > \delta$. So it might be the case that $[z, x'] \notin \mathfrak{D}^{\mathrm{si}}_{\delta,S}$. However, we now have:

$$\begin{aligned}&[z', x'] \in \mathfrak{D}^{\mathrm{si}}_{\delta,S}, \text{ with } z' \text{ as a } \delta\text{-sink, and}\\ &[z, z'] \in \mathfrak{D}^{\mathrm{si}}_{\delta,S}, \text{ with } z \text{ as a } \delta\text{-sink.}\end{aligned}$$

Then we have:

$$\begin{aligned}\partial^{\delta}_1(\gamma) = \partial^{\delta}_1([z, x']) = x' - z &= z' - z + x' - z'\\ &= \partial^{\delta}_1([z, z']) + \partial^{\delta}_1([z', x'])\\ &= \partial^{\delta,S}_1([z, z']) + \partial^{\delta,S}_1([z', x']) \in \mathrm{im}(\partial^{\delta,S}_1),\end{aligned}$$

where the last equality is defined because we have checked that $[z, z'], [z', x'] \in \mathfrak{D}^{\mathrm{si}}_{\delta,S}$. Thus $\mathrm{im}(\partial^{\delta}_1) \subseteq \mathrm{im}(\partial^{\delta,S}_1)$, and the reverse inclusion follows by a similar argument.

Since $\delta \in \mathbb{R}$ was arbitrary, this shows that $\mathrm{im}(\partial^{\delta}_1) = \mathrm{im}(\partial^{\delta,S}_1)$ for each $\delta \in \mathbb{R}$. Previously we had $\ker(\partial^{\delta}_0) = \ker(\partial^{\delta,S}_0)$ for each $\delta \in \mathbb{R}$. It then follows that $H_0(\mathfrak{D}^{\mathrm{si}}_{\delta}) = H_0(\mathfrak{D}^{\mathrm{si}}_{\delta,S})$ for each $\delta \in \mathbb{R}$.

Next let $\delta' \geqslant \delta \in \mathbb{R}$, and for any $k \in \mathbb{Z}_+$, let $f^{\delta,\delta'}_k : C^{\delta}_k \to C^{\delta'}_k, g^{\delta,\delta'}_k : C^{\delta,S}_k \to C^{\delta',S}_k$ denote the chain maps induced by the inclusions $\mathfrak{D}^{\mathrm{si}}_{\delta} \hookrightarrow \mathfrak{D}^{\mathrm{si}}_{\delta'}, \mathfrak{D}^{\mathrm{si}}_{\delta,S} \hookrightarrow \mathfrak{D}^{\mathrm{si}}_{\delta',S}$. Since $\mathfrak{D}^{\mathrm{si}}_{\delta}$ and $\mathfrak{D}^{\mathrm{si}}_{\delta,S}$ have the same 0-simplices at each $\delta \in \mathbb{R}$, we know that $f^{\delta,\delta'}_0 \equiv g^{\delta,\delta'}_0$.

Let $\gamma \in \ker(\partial^{\delta}_0) = \ker(\partial^{\delta,S}_0)$, and let $\gamma + \mathrm{im}(\partial^{\delta}_1) \in H_0(\mathfrak{D}^{\mathrm{si}}_{\delta})$. Then observe that

$$\begin{aligned}(f^{\delta,\delta'}_0)_{\#}(\gamma + \mathrm{im}(\partial^{\delta}_1)) &= f^{\delta,\delta'}_0(\gamma) + \mathrm{im}(\partial^{\delta'}_1) && (f^{\delta,\delta'}_0 \text{ is a chain map})\\ &= g^{\delta,\delta'}_0(\gamma) + \mathrm{im}(\partial^{\delta'}_1) && (f^{\delta,\delta'}_0 \equiv g^{\delta,\delta'}_0)\\ &= g^{\delta,\delta'}_0(\gamma) + \mathrm{im}(\partial^{\delta',S}_1) && (\mathrm{im}(\partial^{\delta'}_1) = \mathrm{im}(\partial^{\delta',S}_1))\\ &= (g^{\delta,\delta'}_0)_{\#}(\gamma + \mathrm{im}(\partial^{\delta,S}_1)). && (g^{\delta,\delta'}_0 \text{ is a chain map})\end{aligned}$$

Thus $(f^{\delta,\delta'}_0)_{\#} = (g^{\delta,\delta'}_0)_{\#}$ for each $\delta' \geqslant \delta \in \mathbb{R}$. Since we also have $H_0(\mathfrak{D}^{\mathrm{si}}_{\delta}) = H_0(\mathfrak{D}^{\mathrm{si}}_{\delta,S})$ for each $\delta \in \mathbb{R}$, we can then apply the Persistence Equivalence Theorem (Theorem 16) to conclude the proof. □

Example 21. Consider the three node dissimilarity networks (X, ω_X) and (Y, ω_Y) in Figure 4. Note that (Y, ω_Y) coincides with $S_X(a, c)$. We present both the Dowker and Rips persistence barcodes obtained from these networks. Note that the Dowker persistence barcode is sensitive to the difference between (X, ω_X) and (Y, ω_Y), whereas the Rips barcode is blind to this difference. We refer the reader to §7 for details on how we compute these barcodes.

To show how the Dowker complex is constructed, we also list the Dowker sink complexes of the networks in Figure 4, and also the corresponding homology dimensions across a range of resolutions. Note that when

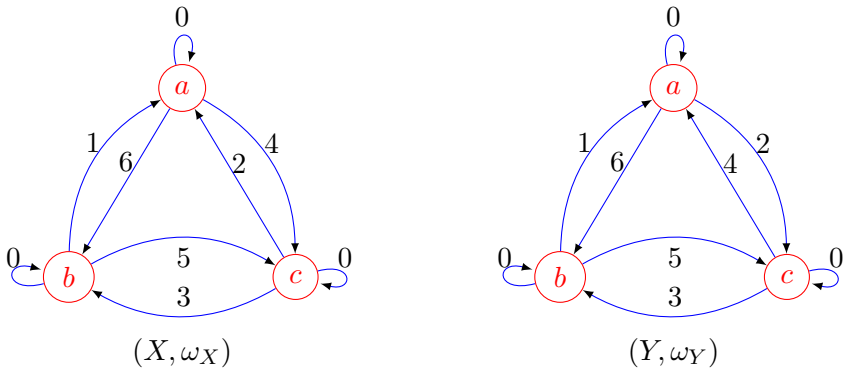

FIGURE 4. (Y, ω_Y) is the (a, c)-swap of (X, ω_X).

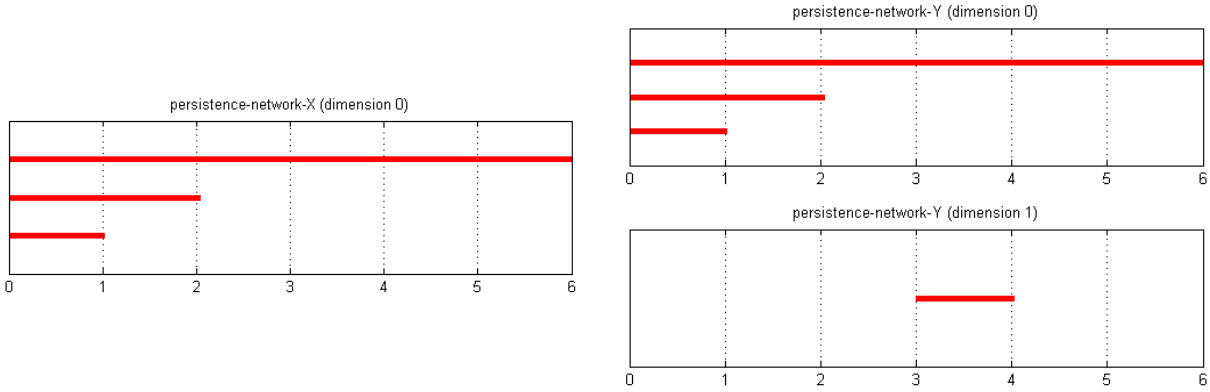
persistence-network-X (dimension 0)
0 1 2 3 4 5 6
persistence-network-Y (dimension 0)
0 1 2 3 4 5 6
persistence-network-Y (dimension 1)
0 1 2 3 4 5 6

FIGURE 5. Dowker persistence barcodes of networks (X, ω_X) and (Y, ω_Y) from Figure 4.

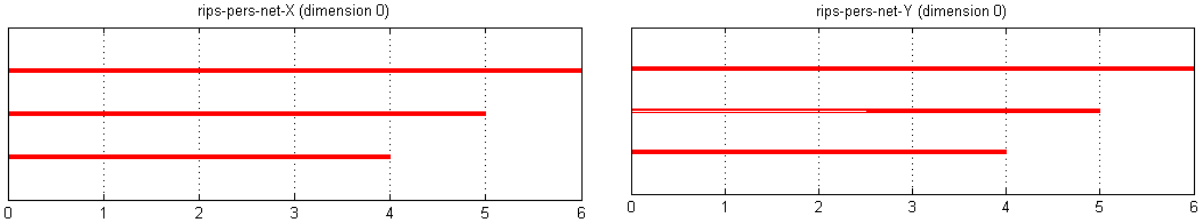

FIGURE 6. Rips persistence barcodes of networks (X, ω_X) and (Y, ω_Y) from Figure 4. Note that the Rips diagrams indicate no persistent homology in dimensions higher than 0, in contrast with the Dowker diagrams in Figure 5.

we write $[a, b](a)$, we mean that a is a sink corresponding to the simplex $[a, b]$.

$$\begin{aligned}
\mathfrak{D}^{\mathrm{si}}_{0,X} &= \{[a], [b], [c]\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{0,X})) &= 0 \\
\mathfrak{D}^{\mathrm{si}}_{1,X} &= \{[a], [b], [c], [a, b](a)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{1,X})) &= 0 \\
\mathfrak{D}^{\mathrm{si}}_{2,X} &= \{[a], [b], [c], [a, b](a), [a, c](a), [b, c](a), [a, b, c](a)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{2,X})) &= 0 \\
\mathfrak{D}^{\mathrm{si}}_{3,X} &= \{[a], [b], [c], [a, b](a), [a, c](a), [b, c](a), [a, b, c](a)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{3,X})) &= 0
\end{aligned}$$

$$\begin{aligned}
\mathfrak{D}^{\mathrm{si}}_{0,Y} &= \{[a],[b],[c]\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{0,Y})) &= 0\\
\mathfrak{D}^{\mathrm{si}}_{1,Y} &= \{[a],[b],[c],[a,b](a)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{1,Y})) &= 0\\
\mathfrak{D}^{\mathrm{si}}_{2,Y} &= \{[a],[b],[c],[a,b](a),[a,c](c)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{2,Y})) &= 0\\
\mathfrak{D}^{\mathrm{si}}_{3,Y} &= \{[a],[b],[c],[a,b](a),[a,c](c),[b,c](b)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{3,Y})) &= 1\\
\mathfrak{D}^{\mathrm{si}}_{4,Y} &= \{[a],[b],[c],[a,b](a),[a,c](a),[b,c](a),[a,b,c](a)\} & \dim(H_1(\mathfrak{D}^{\mathrm{si}}_{4,Y})) &= 0
\end{aligned}$$

Note that for $\delta \in [3,4)$, $\dim(H_1(\mathfrak{D}^{\mathrm{si}}_{\delta,Y})) = 1$, whereas $\dim(H_1(\mathfrak{D}^{\mathrm{si}}_{\delta,X})) = 0$ for each $\delta \in \mathbb{R}$.

Based on the discussion in Remark 19, Proposition 20, and Example 21, we conclude the following:

Moral: *Unlike Rips persistence diagrams, Dowker persistence diagrams are truly sensitive to asymmetry.*

Proceeding beyond Example 21, we now provide a family of asymmetric networks for which Dowker persistence captures more relevant information than Rips persistence.

5.2.1. *Cycle networks.* For each $n \in \mathbb{N}$, let (X_n, E_n, W_{E_n}) denote the weighted graph with vertex set $X_n := \{x_1, x_2, \ldots, x_n\}$, edge set $E_n := \{(x_1,x_2),(x_2,x_3),\ldots,(x_{n-1},x_n),(x_n,x_1)\}$, and edge weights $W_{E_n} : E_n \to \mathbb{R}$ given by writing $W_{E_n}(e) = 1$ for each $e \in E_n$. Next let $\omega_{G_n} : X_n \times X_n \to \mathbb{R}$ denote the shortest path distance induced on $X_n \times X_n$ by W_{E_n}. Then we write $G_n := (X_n, \omega_{G_n})$ to denote the network with node set X_n and weights given by ω_{G_n}. Note that $\omega_{G_n}(x,x) = 0$ for each $x \in X_n$.

We say that G_n is the *cycle network of length n*. One can interpret cycle networks as being highly asymmetric, because for every consecutive pair of nodes (x_i, x_{i+1}) in a graph G_n, where $1 \leqslant i \mod (n) \leqslant n$, we have $\omega_{G_n}(x_i, x_{i+1}) = 1$, whereas $\omega_{G_n}(x_{i+1}, x_i) = \mathrm{diam}(G_n) = n-1$, which is much larger than 1 when n is large.

To provide further evidence that Dowker persistence is sensitive to asymmetry, we computed both the Rips and Dowker persistence diagrams, in dimensions 0 and 1, of cycle networks G_n, for values of n between 3 and 6. Computations were carried out using `Javaplex` in Matlab with $\mathbb{Z}_2$ coefficients. The results are presented in Figure 7. Based on our computations, we were able to conjecture and prove the result in Theorem 22, which gives a precise characterization of the 1-dimensional Dowker persistence diagram of a cycle network G_n, for any n. Furthermore, the 1-dimensional Dowker persistence barcode for any G_n contains only one persistent interval, which agrees with our intuition that there is only one nontrivial loop in G_n. On the other hand, for large n, the 1-dimensional Rips persistence barcodes contain more than one persistent interval. This can be seen in the Rips persistence barcode of G_6, presented in Figure 7. Moreover, for $n = 3, 4$, the 1-dimensional Rips persistence barcode does not contain any persistent interval at all. This suggests that Dowker persistence diagrams/barcodes are an appropriate method for analyzing cycle networks, and perhaps asymmetric networks in general.

Notation. In the remainder of this section, we will prove results involving Dowker sink complexes of the cycle networks G_n and associated vector spaces at a range of resolutions δ. For convenience, we will write $\mathfrak{D}^{\mathrm{si}}_{\delta} := \mathfrak{D}^{\mathrm{si}}_{\delta,G_n}$ (where n will be fixed) and $C^{\delta}_k := C_k(\mathfrak{D}^{\mathrm{si}}_{\delta})$, the k-chain vector space associated to $\mathfrak{D}^{\mathrm{si}}_{\delta}$ for each $k \in \mathbb{Z}_+$. For each $k \in \mathbb{Z}_+$, the boundary map from C^{δ}_k to C^{δ}_{k-1} will be denoted ∂^{δ}_k. Whenever we write x_i to denote a vertex of G_n, the subscript i should be understood as $i \pmod n$. We write e_i to denote the 1-simplex $[x_i, x_{i+1}]$ for each $1 \leqslant i \leqslant n$, where x_{n+1} is understood to be x_1. Given an element $\gamma \in \ker(\partial^{\delta}_k) \subseteq C^{\delta}_1$, we will write $\langle\gamma\rangle_{\delta}$ to denote its equivalence class in the quotient vector space $\ker(\partial^{\delta}_k)/\mathrm{im}(\partial^{\delta}_k)$. We will refer to the operation of taking this quotient as *passing to homology*.

The following theorem contains the characterization result for 1-dimensional Dowker persistence diagrams of cycle networks.

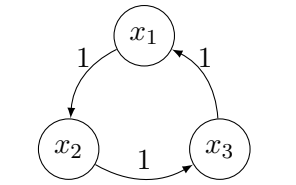

The cycle network G_3

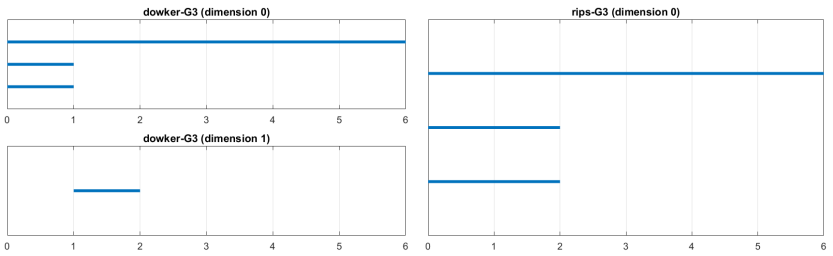

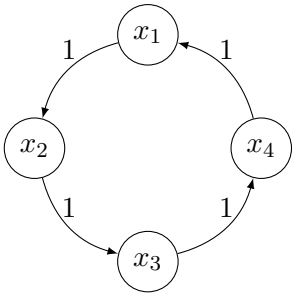

The cycle network G_4

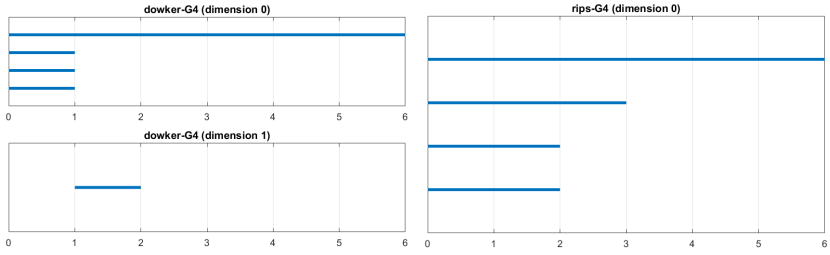

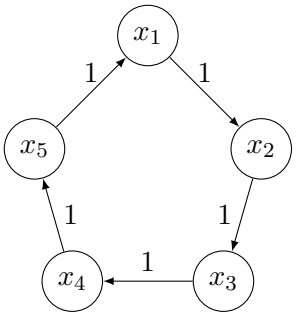

The cycle network G_5

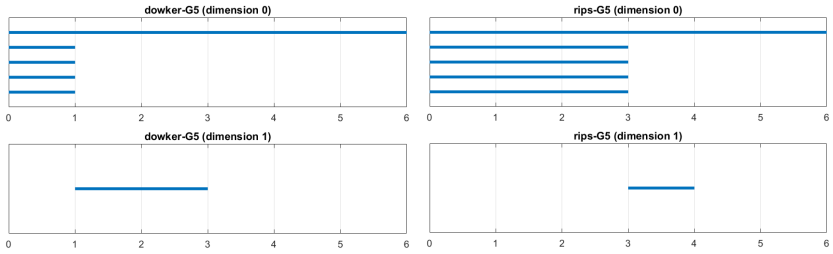

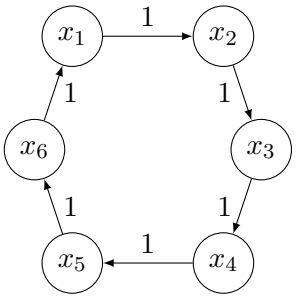

The cycle network G_6

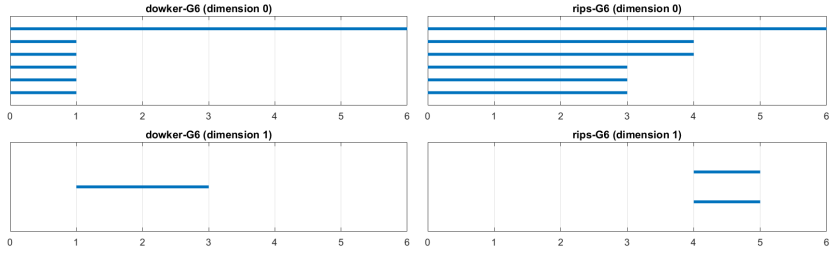

FIGURE 7. The first column contains illustrations of cycle networks G_3, G_4, G_5 and G_6. The second column contains the corresponding Dowker persistence barcodes, in dimensions 0 and 1. Note that the persistent intervals in the 1-dimensional barcodes agree with the result in Theorem 22. The third column contains the Rips persistence barcodes of each of the cycle networks. Note that for $n = 3, 4$, there are no persistent intervals in dimension 1. On the other hand, for $n = 6$, there are two persistent intervals in dimension 1.

Theorem 22. *Let $G_n = (X_n, \omega_{G_n})$ be a cycle network for some $n \in \mathbb{N}$. Then we obtain:*

$$\mathrm{Dgm}_1^{\mathrm{si}}(G_n) = \left\{(1, \lceil n/2 \rceil) \in \mathbb{R}^2\right\}.$$

Thus $\mathrm{Dgm}_1^{\mathrm{si}}(G_n)$ consists of precisely the point $(1, \lceil n/2 \rceil) \in \mathbb{R}^2$ with multiplicity 1.

Lemma 23. *Let $\delta \geqslant 0$ and let $G_n = (X_n, \omega_{G_n})$ be a cycle network, for any natural number $n \geqslant 3$. Consider the boundary map $\partial_2^\delta : C_2^\delta \to C_1^\delta$. Then $\mathrm{im}(\partial_2^\delta)$ is generated by 1-chains of the form $\partial_2^\delta([x_1, x_i, x_j])$, where $1 < i < j \leqslant n$.*

Proof of Lemma 23. Let $\partial_2^\delta([x_i, x_j, x_k]) \in \mathrm{im}\, \partial_2^\delta$, where $i < j < k$. If $n = 3$ or if $i = 1$, we are done. Suppose $n \geqslant 4$ and $i > 1$. Note that:

$$\partial_3^\delta([x_1, x_i, x_j, x_k]) = [x_i, x_j, x_k] - [x_1, x_j, x_k] + [x_1, x_i, x_k] - [x_1, x_i, x_j].$$

Since $\partial_2^\delta \circ \partial_3^\delta = 0$, we have:

$$\partial_2^\delta([x_i, x_j, x_k]) = \partial_2^\delta([x_1, x_j, x_k]) - \partial_2^\delta([x_1, x_i, x_k]) + \partial_2^\delta([x_1, x_i, x_j]).$$

Since $\partial_2^\delta([x_i, x_j, x_k]) \in \mathrm{im}\, \partial_2^\delta$ was arbitrary, we are done. □

Proof of Theorem 22. The proof occurs in three stages: first we show that a 1-cycle appears at $\delta = 1$, next we show that this 1-cycle does not become a boundary until $\delta = \lceil n/2 \rceil$, and finally that any other 1-cycle belongs to the same equivalence class upon passing to homology (this shows that the single point in the persistence diagram has multiplicity 1).

First we deal with the easy cases $n = 1, 2$. In either of these cases, $\lceil n/2 \rceil = 1$, so the claim is that $\mathrm{Dgm}_1^{\mathrm{si}}(G_n) = \{(1, 1)\}$, i.e. that the 1-dimensional Dowker persistence diagram is trivial, in the sense of Remark 1. Note that if $n = 1$, there are no 1-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$ for any $\delta \geqslant 0$, and that if $n = 2$, the only 1-simplex in $\mathfrak{D}_\delta^{\mathrm{si}}$ is $[x_1, x_2]$. In either case, $H_1(\mathfrak{D}_\delta^{\mathrm{si}})$ is the trivial vector space, for any $\delta \geqslant 0$. This proves the claim for $n = 1, 2$. In the sequel, we assume $n \geqslant 3$.

Note that for $\delta < 1$, there are no 1-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$, and so $H_1(\mathfrak{D}_\delta^{\mathrm{si}})$ is trivial. Suppose $1 \leqslant \delta < 2$.

Claim 3. There are no 2-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$ for $1 \leqslant \delta < 2$.

Proof. To see this, let x_i, x_j, x_k be any three distinct vertices in X_n. Assume towards a contradiction that there exists $x \in X_n$ such that $(x_i, x), (x_j, x), (x_k, x) \in R_{\delta, X_n}$, where R_{δ, X_n} is as given by Equation 2. Thus $\omega_{G_n}(x_i, x) \in \{0, 1\}$, so either $x = x_i$ or $x = x_{i+1}$. Similarly we get that $x = x_j$ or $x = x_{j+1}$, and that $x = x_k$ or $x = x_{k+1}$. But this is a contradiction, since x_i, x_j, x_k are all distinct. ■

By the claim, there are no 2-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$, so $\mathrm{im}(\partial_2^\delta)$ is trivial and the only 1-chains are linear combinations of e_i terms. Next, we define:

$$v_n := e_1 + e_2 + \ldots + e_n = [x_1, x_2] + [x_2, x_3] + \ldots + [x_n, x_1].$$

Note that $v_n \in C_1^\delta$ for all $\delta \geqslant 1$. One can further verify that $\partial_1^\delta(v_n) = 0$, for any $\delta \geqslant 1$. In other words, v_n is a 1-cycle for any $\delta \geqslant 1$.

Claim 4. Let $1 \leqslant \delta < 2$. Then v_n generates $\ker(\partial_1^\delta) \subseteq C_1^\delta$.

Proof. The only 1-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$ are of the form e_i, for $1 \leqslant i \leqslant n$. So it suffices to show that any linear combination of the e_i terms is a multiple of v_n. Let $u = \sum_{i=1}^n a_i e_i \in \ker(\partial_1^\delta)$, for some $a_1, \ldots, a_n \in \mathbb{K}$. Then,

$$\begin{aligned} 0 = \partial_1^\delta(u) = \sum_{i=1}^n a_i \partial_1^\delta(e_i) &= \sum_{i=1}^n a_i([x_{i+1}] - [x_i]) \\ &= \sum_{i=1}^n (a_{i-1} - a_i)[x_i], \qquad \text{where } x_0 \text{ is understood to be } x_n. \end{aligned}$$

Since all the $[x_i]$ are linearly independent, it follows that $a_1 = a_2 = \ldots = a_n$. Thus it follows that u is a constant multiple of v_n. This proves the claim. ■

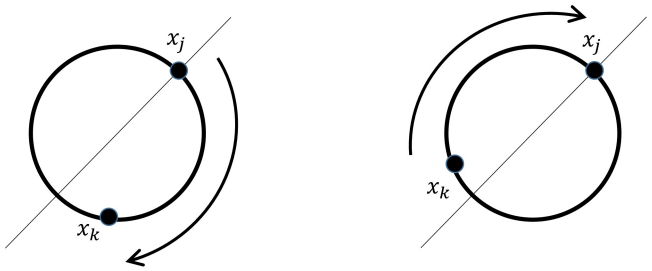

FIGURE 8. Given two points $x_j, x_k \in X_n$, we have either $\omega_{G_n}(x_j, x_k) \leqslant n/2$, or $\omega_{G_n}(x_k, x_j) < n/2$. To see this, note that $\omega_{G_n}(x, x') + \omega_{G_n}(x', x) = n$ for any $x \neq x' \in X_n$.

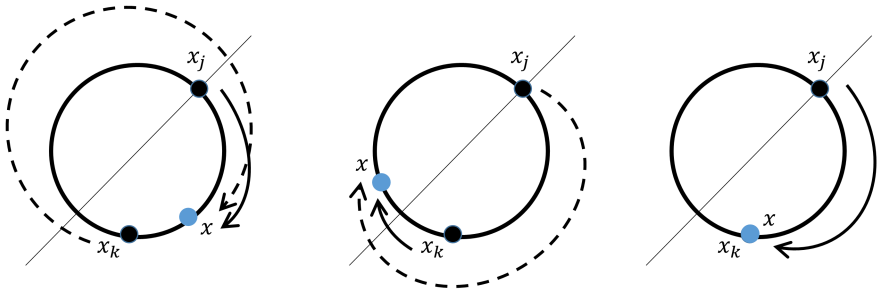

FIGURE 9. Three possible locations for a δ-sink x of a simplex $[x_j, x_k]$, assuming that $\omega_{G_n}(x_j, x_k) \leqslant n/2$. For the figure on the left, note that $\omega_{G_n}(x_k, x) \geqslant n/2 \geqslant \omega_{G_n}(x_j, x_k)$. For the figure in the middle, note that $\omega_{G_n}(x_j, x) \geqslant \omega_{G_n}(x_j, x_k)$. Finally, for the figure on the right, where $x = x_k$, note that $\omega_{G_n}(x_j, x) = \omega_{G_n}(x_j, x_k)$ and $\omega_{G_n}(x_k, x) = 0$.

By the two preceding claims, it follows that $\{\langle v_n \rangle_\delta\}$ is a basis for $H_1(\mathfrak{D}_\delta^{\mathrm{si}})$, for $\delta \in [1, 2)$. More specifically, $\langle v_n \rangle_\delta$ is a cycle that appears at $\delta = 1$ and does not become a boundary until at least $\delta = 2$, and any other cycle in C_1^δ, for $\delta \in [1, 2)$, is in the linear span of v_n. Next, suppose $\delta \geqslant 2$. Note that this allows the appearance of cycles that are not in the span of v_n. In the next claim, we show that upon passing to homology, the equivalence class of any such cycle coincides with that of v_n. This will show that there can be at most one nontrivial element in $\mathrm{Dgm}_1^{\mathrm{si}}(G_n)$.

Claim 5. Let $\delta \geqslant 2$, and let $y = \sum_{i=1}^p a_i \sigma_i \in \ker(\partial_1^\delta)$ for some $p \in \mathbb{N}$, some $a_1, \ldots, a_p \in \mathbb{K}$, and some $\sigma_1, \ldots, \sigma_p \in \mathfrak{D}_\delta^{\mathrm{si}}$. Then there exists a choice of coefficients $(b_i)_{i=1}^n \in \mathbb{K}^n$ such that $z = \sum_{i=1}^n b_i e_i \in \ker(\partial_1^\delta)$ and $y - z \in \mathrm{im}(\partial_2^\delta)$. Moreover, we obtain $\langle y \rangle_\delta = \langle z \rangle_\delta = \langle v_n \rangle_\delta$ upon passing to homology.

Proof. To see this, fix $\sigma_i \in \mathfrak{D}_\delta^{\mathrm{si}}$, and write $\sigma_i = [x_j, x_k]$ for some $1 \leqslant j, k \leqslant n$. If $k = j + 1$ (resp. $k = j - 1$), then we already have $\sigma_i = e_j$ (resp. $\sigma_i = e_k$), so there is nothing more to show. Assume $k \notin \{j + 1, j - 1\}$. Since $\omega_{G_n}(x_j, x_k) + \omega_{G_n}(x_k, x_j) = n$, we have two cases: (1) $\omega_{G_n}(x_j, x_k) \leqslant n/2$, or (2) $\omega_{G_n}(x_k, x_j) < n/2$. In the first case, we have $k = j + l$ for some integer $l \in [2, n/2]$ (all numbers are taken modulo n). In the second case, $j = k + l$ for some integer $l \in [2, n/2)$ (also modulo n). The situation is illustrated in Figure 8. Both cases are similar, so we only prove the case $\omega_{G_n}(x_j, x_k) \leqslant n/2$.

Recall that any δ-sink $x \in X_n$ for $[x_j, x_k]$ satisfies $\max(\omega_{G_n}(x_j, x), \omega_{G_n}(x_k, x)) \leqslant \delta$, by the δ-sink condition (Equation 3). Also note that such a δ-sink x satisfies

$$\max(\omega_{G_n}(x_j, x), \omega_{G_n}(x_k, x)) \geqslant \omega_{G_n}(x_j, x_k),$$

as can be seen from Figure 9. So whenever some $x \in X_n$ is a δ-sink for $[x_j, x_k]$, we have x_k as a valid δ-sink for $[x_j, x_k]$. Since $[x_j, x_k] \in \mathfrak{D}_\delta^{\mathrm{si}}$, it must have a δ-sink $x \in X_n$. Thus x_k is a valid δ-sink for $[x_j, x_k]$.

Next let $l \in [2, n/2]$ be an integer such that $k = j + l$ (modulo n). Notice that:

$$0 = \omega_{G_n}(x_k, x_k) = \omega_{G_n}(x_{j+l}, x_k) < \omega_{G_n}(x_{j+l-1}, x_k) < \ldots < \omega_{G_n}(x_{j+1}, x_k) < \omega_{G_n}(x_j, x_k) \leqslant \delta.$$

Then observe that:

$$[x_j, x_{j+1}, x_k], [x_{j+1}, x_{j+2}, x_k], \ldots, [x_{k-2}, x_{k-1}, x_k] \in \mathfrak{D}_\delta^{\mathrm{si}},$$

since x_k is a δ-sink for all these 2-simplices. One can then verify the following:

$$\begin{aligned}
&\partial_2^\delta \left([x_j, x_{j+1}, x_k] + [x_{j+1}, x_{j+2}, x_k] + \ldots + [x_{k-2}, x_{k-1}, x_k] \right) \\
&= \partial_2^\delta \left(\sum_{q=0}^{k-j-2} [x_{j+q}, x_{j+q+1}, x_k] \right) \\
&= \sum_{q=0}^{k-j-2} [x_{j+q+1}, x_k] - \sum_{q=0}^{k-j-2} [x_{j+q}, x_k] + \sum_{q=0}^{k-j-2} [x_{j+q}, x_{j+q+1}] \\
&= \sum_{q=0}^{k-j-2} [x_{j+q+1}, x_k] - [x_j, x_k] - \sum_{q=0}^{k-j-3} [x_{j+q+1}, x_k] + \sum_{q=0}^{k-j-2} [x_{j+q}, x_{j+q+1}] \\
&= [x_j, x_{j+1}] + [x_{j+1}, x_{j+2}] + \ldots + [x_{k-1}, x_k] - [x_j, x_k] \\
&= e_j + e_{j+1} + \ldots + e_{k-1} - \sigma_i.
\end{aligned}$$

Thus $a_i(e_j + e_{j+1} + \ldots + e_{k-1}) - a_i\sigma_i \in \mathrm{im}(\partial_2^\delta)$. Repeating this process for all σ_i, $i \in \{1, \ldots, p\}$, we may obtain the coefficients $(b_i)_{i=1}^n$ such that $\sum_{i=1}^p a_i\sigma_i - \sum_{i=1}^n b_i e_i \in \mathrm{im}(\partial_2^\delta)$. Let $z = \sum_{i=1}^n b_i e_i$. Then $y - z \in \mathrm{im}(\partial_2^\delta)$. Moreover, since $\partial_1^\delta \circ \partial_2^\delta = 0$, it follows that $\partial_1^\delta(y) - \partial_1^\delta(z) = 0$, so $z \in \ker(\partial_1^\delta)$.

Finally, note that an argument analogous to that of Claim 4 shows that $b_1 = b_2 = \ldots = b_n$. Hence it follows that z is a multiple of v_n. Thus $\langle z \rangle_\delta = \langle v_n \rangle_\delta$. This proves the claim. ■

By Claims 4 and 5, it follows that $H_1(\mathfrak{D}_\delta^{\mathrm{si}})$ is generated by $\langle v_n \rangle_\delta$ for all $\delta \geqslant 1$, so $\dim(H_1(\mathfrak{D}_\delta^{\mathrm{si}})) \leqslant 1$ for all $\delta \geqslant 1$. It remains to show that $\langle v_n \rangle_\delta$ does not become trivial until $\delta = \lceil n/2 \rceil$.

The cases $n = 3, 4$ can now be completed quickly, so we focus on these simpler situations first. For either of $n = 3, 4$, we have $\lceil n/2 \rceil = 2$. Suppose $\delta = 2$ and $n = 3$. Then we have $[x_1, x_2, x_3] \in \mathfrak{D}_\delta^{\mathrm{si}}$ because $\mathrm{diam}(G_n) = 2$ and any of x_1, x_2, x_3 can be a 2-sink for $[x_1, x_2, x_3]$. Then,

$$\partial_2^\delta([x_1, x_2, x_3]) = [x_2, x_3] - [x_1, x_3] + [x_1, x_2] = e_1 + e_2 + e_3 = v_3.$$

Recall that by Claim 3, $v_3 \notin \mathrm{im}(\partial_2^\delta)$ for any $\delta < 2$. Thus by Claim 4 and the preceding equation, v_3 generates $\ker(\partial_1^\delta)$ for $1 \leqslant \delta < 2$, and becomes a boundary for precisely $\delta \geqslant 2$. Thus $\mathrm{Dgm}_1^{\mathrm{si}}(G_3) = \{(1, 2)\}$. Next, suppose $\delta = 2$ and $n = 4$. Then we have $[x_1, x_2, x_3], [x_1, x_3, x_4] \in \mathfrak{D}_\delta^{\mathrm{si}}$ with x_3, x_1 as 2-sinks, respectively. By a direct computation, we then have:

$$\partial_2^\delta([x_1, x_2, x_3] + [x_1, x_3, x_4]) = e_1 + e_2 + e_3 + e_4 = v_4.$$

By following the same argument as for the case $n = 3$, we see that $\mathrm{Dgm}_1^{\mathrm{si}}(G_4) = \{(1, 2)\}$.

In the sequel, we assume that $n > 4$. Recall that it remains to show that $\langle v_n \rangle_\delta$ does not become trivial until $\delta = \lceil n/2 \rceil$, and that $\langle v_n \rangle_\delta = 0$ for all $\delta \geqslant \lceil n/2 \rceil$. We have already shown that $\langle v_n \rangle_\delta$ is not trivial for $\delta \in [1, 2)$. We proceed by defining the following:

$$\gamma_n := [x_1, x_2, x_3] + [x_1, x_3, x_4] + \ldots + [x_1, x_{n-1}, x_n] = \sum_{i=1}^{n-2} [x_1, x_{i+1}, x_{i+2}].$$

Claim 6. For each $\delta \geqslant \lceil n/2 \rceil$, we have $\gamma_n \in C_2^\delta$ and $\partial_2^\delta(\gamma_n) = v_n$. In particular, $\langle v_n \rangle_\delta = 0$ for all such δ.

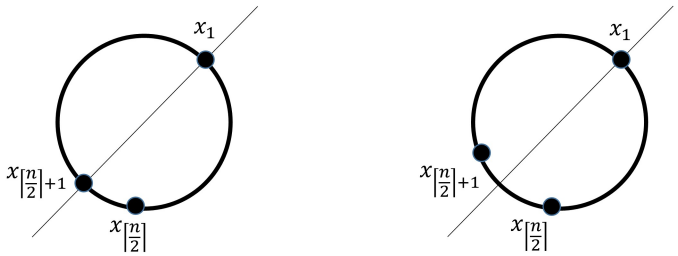

FIGURE 10. Placement of $x_{\lceil n/2 \rceil}$ and $x_{\lceil n/2 \rceil - 1}$, depending on whether n is even or not.

Proof. Let $\delta \geqslant \lceil n/2 \rceil$. Notice that

$$\omega_{G_n}(x_{\lceil n/2 \rceil + 1}, x_1) = n - \omega_{G_n}(x_1, x_{\lceil n/2 \rceil + 1}) = n - \lceil n/2 \rceil \leqslant n/2 \leqslant \lceil n/2 \rceil \leqslant \delta,$$

so $\omega_{G_n}(x_i, x_1) \leqslant \delta$ for each $i \in \{\lceil n/2 \rceil + 1, \lceil n/2 \rceil + 2, \ldots, n\}$. Then for each $i \in \{\lceil n/2 \rceil + 1, \lceil n/2 \rceil + 2, \ldots, n-1\}$, we have $[x_i, x_{i+1}, x_1] \in \mathfrak{D}_\delta^{\mathrm{si}}$, with x_1 as a δ-sink.

Also notice that for each $i \in \{1, \ldots, \lceil n/2 \rceil\}$,

$$\omega_{G_n}(x_i, x_{\lceil n/2 \rceil + 1}) \leqslant \omega_{G_n}(x_1, x_{\lceil n/2 \rceil + 1}) = \lceil n/2 \rceil \leqslant \delta,$$

so $\omega_{G_n}(x_i, x_{\lceil n/2 \rceil + 1}) \leqslant \delta$. Thus for any $i \in \{2, \ldots, \lceil n/2 \rceil\}$, we have $[x_1, x_i, x_{i+1}] \in \mathfrak{D}_\delta^{\mathrm{si}}$, with $x_{\lceil n/2 \rceil + 1}$ as a δ-sink.

Combining the two preceding observations, we see that for any $i \in \{2, \ldots, n-2\}$, we have $[x_1, x_{i+1}, x_{i+2}] \in \mathfrak{D}_\delta^{\mathrm{si}}$. It follows that $\gamma_n \in C_2^\delta$.

Next we observe the following:

$$\begin{aligned}
\partial_2^\delta(\gamma_n) &= \partial_2^\delta \left(\sum_{i=1}^{n-2} [x_1, x_{i+1}, x_{i+2}] \right) \\
&= \sum_{i=1}^{n-2} [x_{i+1}, x_{i+2}] - \sum_{i=1}^{n-2} [x_1, x_{i+2}] + \sum_{i=1}^{n-2} [x_1, x_{i+1}] \\
&= \sum_{i=1}^{n-2} [x_{i+1}, x_{i+2}] - \sum_{i=1}^{n-2} [x_1, x_{i+2}] + [x_1, x_2] + \sum_{i=2}^{n-2} [x_1, x_{i+1}] \\
&= \sum_{i=1}^{n-2} [x_{i+1}, x_{i+2}] + [x_1, x_2] - [x_1, x_n] = v_n.
\end{aligned}$$

It follows that for any $\delta \geqslant \lceil n/2 \rceil$, we have $v_n \in \mathrm{im}(\partial_2^\delta)$, and so $\langle v_n \rangle_\delta = 0$ for each such δ. ∎

Claim 7. There does not exist $\delta \in [2, \lceil n/2 \rceil)$ such that $\langle v_n \rangle_\delta$ is trivial.

Proof. Let $2 \leqslant \delta < \lceil n/2 \rceil$. As a first step, we wish to show that $\gamma_n \notin C_2^\delta$. For this step, it suffices to show that the 2-simplex $\sigma := [x_1, x_{\lceil n/2 \rceil}, x_{\lceil n/2 \rceil + 1}]$ does not belong to $\mathfrak{D}_\delta^{\mathrm{si}}$. The placement of $x_{\lceil n/2 \rceil}$ and $x_{\lceil n/2 \rceil + 1}$ is illustrated in Figure 10.

By an argument similar to that used in Figure 9, one can verify that there exists a δ-sink for σ if and only if at least one of $x_1, x_{\lceil n/2 \rceil}, x_{\lceil n/2 \rceil + 1}$ is a δ-sink for σ. But note the following:

$$\begin{aligned}
\omega_{G_n}(x_{\lceil n/2 \rceil}, x_1) = n - (\lceil n/2 \rceil - 1) &= \begin{cases} n/2 + 1 & : n \text{ even} \\ \lceil n/2 \rceil & : n \text{ odd} \end{cases} \\
&\geqslant \lceil n/2 \rceil > \delta,
\end{aligned}$$

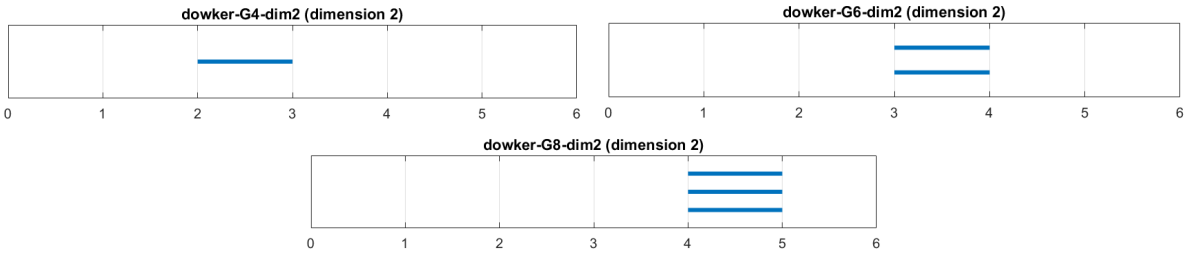

FIGURE 11. Sample 2-dimensional Dowker persistence barcodes for cycle networks G_4, G_6, G_8. In our experiments, 2-dimensional Dowker persistence barcodes for G_n were always empty for n odd.

so x_1 cannot be a δ-sink for σ. Similarly we note that $\omega_{G_n}(x_{\lceil n/2\rceil+1}, x_{\lceil n/2\rceil}) = n > \delta$ and $\omega_{G_n}(x_1, x_{\lceil n/2\rceil+1}) = \lceil n/2\rceil > \delta$, so neither $x_{\lceil n/2\rceil}$ nor $x_{\lceil n/2\rceil+1}$ can be δ-sinks for σ. Thus $\sigma \notin \mathfrak{D}_\delta^{\mathrm{si}}$, and so $\gamma_n \notin C_2^\delta$.

Suppose there exists $\gamma' \in C_2^\delta$ such that $\partial_2^\delta(\gamma') = v_n$. Then by Lemma 23, we may write $\partial_2^\delta(\gamma') = v_n$ as a sum of the form $\sum_{i,j} a_{ij}\partial_2^\delta([x_1, x_i, x_j])$, where each $a_{ij} \in \mathbb{K}$ and $1 < i < j \leqslant n$. But then, by a linear independence argument analogous to that used in Claim 4, we must have $\gamma_n = \gamma'$, a contradiction. Thus $v_n \notin \mathrm{im}(\partial_2^\delta)$. It follows that $\langle v_n \rangle_\delta \neq 0$. ■

Thus we have shown that v_n is a nontrivial cycle that appears at $\delta = 1$, and becomes a boundary at exactly $\delta = \lceil n/2\rceil$. Furthermore, we have shown that upon passing to homology, the equivalence classes of all cycles coincide with that of v_n. Thus there is only one off-diagonal point $(1, \lceil n/2\rceil)$ on the 1-dimensional persistence diagram, which appears with multiplicity one. This concludes the proof. □

Remark 24. From our experimental results (see Figure 7), it seems that the 1-dimensional Rips persistence diagram of a cycle network does not admit a characterization as simple as that of Theorem 22. A characterization of the Rips persistence diagrams of a cycle network is left as future work. We remark that the authors of [AA15] studied a construction called a *directed cyclic graph* which is slightly more general than a cycle network, and characterized the Rips persistence diagram of its underlying *undirected* graph. Moreover, they gave results for a topological invariant called *homotopy type*, which is stronger than homology. From our perspective, it seems that a full characterization of the Rips persistence diagram of a cycle network would likely be possible, and would be a useful example in the study of persistence of directed networks.

Remark 25. Our motivation for studying cycle networks is that they are a directed analogue of circles, and we were interested in seeing if the 1-dimensional Dowker persistence diagram would be able to capture this analogy. Theorem 22 shows that this is indeed the case: we get a single nontrivial 1-dimensional persistence interval, which is what we would expect when computing the persistent homology of a circle in the metric space setting. We further studied the 2-dimensional Dowker persistence diagrams of cycle networks. Our computational examples, some of which are illustrated in Figure 11, enable us to make the following conjecture:

Conjecture 1. Let $n \in \mathbb{N}$, and let G_n be a cycle network. If n is odd, then $\mathrm{Dgm}_2^{\mathrm{si}}(G_n)$ is trivial. If n is even, then $\mathrm{Dgm}_2^{\mathrm{si}}(G_n) = [(\frac{n}{2}, \frac{n}{2}+1) \in \mathbb{R}^2]$, and the multiplicity of this point is $\frac{n}{2} - 1$.

Proving this conjecture, and higher dimensional analogues, is left as further work.

6. RIPS AND DOWKER HIERARCHICAL CLUSTERING METHODS

In this section, we delve deeper into the "meaning" of a persistence barcode, specifically a 0-dimensional Rips or Dowker barcode, and explore some connections to hierarchical clustering methods. Some of the theoretical underpinnings of hierarchical clustering methods in the setting of metric spaces are described in [CM10, CM13]. Hierarchical clustering methods on dissimilarity networks are described in [CMRS13]. To

facilitate the connection to existing literature, we restrict ourselves to dissimilarity networks (see Remark 12) in this section. We will use notation of the form (X, A_X) (as opposed to (X, ω_X)) to signal that we are referring to dissimilarity networks.

As a motivating example, suppose we are given a 0-dimensional Rips persistence barcode obtained from a network (X, A_X). The right endpoint of each persistence interval in this barcode corresponds to a resolution δ at which a 0-cycle becomes a 0-boundary. One can then verify the following implication: a 1-simplex connecting two separated components enters the Rips complex of X at resolution δ. Note that because we have restricted our attention to dissimilarity networks, we know that: (1) $\mathfrak{R}_X^0$ consists of $\mathrm{card}(X)$ singleton components, and (2) $\mathfrak{R}_X^{\delta_F}$ consists of a single connected component, where $\delta_F := \max_{x,x' \in X} A_X(x, x')$. At any $\delta \in (0, \delta_F)$, $\mathfrak{R}_X^{\delta}$ is a union of connected components, and as δ increases, components only join together, but do not separate.

The description in the preceding paragraph also applies to the Dowker complex of (X, A_X). We remark that this hierarchical description of connected components can be realized as a formal structure called a *dendrogram*. We will momentarily give a precise definition of a dendrogram, but we first note two reasons for studying a dendrogram in addition to a 0-dimensional persistence barcode: (1) a dendrogram has a visual representation as a rooted tree that is just as easy to study as a persistence barcode, and (2) a dendrogram keeps track of the resolutions at which components join together *in addition to* the labels of the components that join together, whereas a 0-dimensional persistence barcode forgets the latter. In the sequel, we will define dendrograms in general, and then the dendrograms that arise from the Rips and Dowker filtrations of a network.

Recall that given a finite set X a *partition* of X is any collection $P = \{B_1, \ldots, B_k\}$ where each B_i is a subset of X, often referred to as a *block* or cluster of the partition P. Different blocks of P also need to be disjoint: for $i \neq j$ one has $B_i \cap B_j = \varnothing$, and the totality of the blocks must cover X completely: $\cup_{i=1}^k B_k = X$. From now on, for a finite set X we denote by $\mathrm{Part}(X)$ the set of all partitions of X.

Definition 3. Let X be a finite set. A *dendrogram* over X is a function $D_X : \mathbb{R}_+ \to \mathrm{Part}(X)$ such that:

(1) (hierarchy) For $t' \geqslant t$, $D_X(t)$ refines $D_X(t')$, meaning that every block of $D_X(t)$ is contained in a block of $D_X(t')$.
(2) $\exists t_F \in \mathbb{R}_+$ such that for all $t \geqslant t_F$, $D_X(t) = \{X\}$, i.e. the one block partition of X.
(3) $D_X(0)$ consists of singletons $\{x\}$ for $x \in X$, i.e. is the finest possible partition of X.
(4) (right continuity) For all $t \in \mathbb{R}_+$ there exists $\varepsilon > 0$ such that $D_X(t') = D_X(t)$ for all $t' \in [t, t + \varepsilon]$.

The parameter t above is referred to as *resolution*. Conditions 2 and 3 are called *boundary conditions*, and they specify the resolutions at which all the nodes of X are clustered together, and at which we have the only clusters are the singletons. Condition 1 (*hierarchy*) emphasizes that as the resolution t increases, clusters can only be combined, not separated. The final condition, *right continuity*, is included for technical reasons.

Any function assigning a dendrogram to a network is called a *network hierarchical clustering method (nHCM)*. To see how an nHCM can be defined, first observe that a partition P on a finite set X is naturally equivalent to an equivalence relation $\sim_P$, as follows: for any $x, x' \in X$, we write $x \sim_P x'$ if and only if x, x' belong to the same block of P. So one method for defining an nHCM would be to specify a procedure that assigns to any network (X, ω_X) a one-parameter family of equivalence relations $\{\sim_\delta\}_{\delta \in \mathbb{R}_+}$, such that the associated partitions of X satisfy the conditions for a dendrogram. More specifically, the family $\{\sim_\delta\}_{\delta \in \mathbb{R}_+}$ is required to satisfy the following conditions:

(1) If $x \sim_\delta x'$, then $x \sim_{\delta'} x'$, whenever $\delta' \geqslant \delta \in \mathbb{R}_+$.
(2) There exists $\delta_F \in \mathbb{R}_+$ such that for all $\delta \geqslant \delta_F$, we have $x \sim_\delta x'$ for any $x, x' \in X$.
(3) $x \sim_0 x'$ if and only if $x = x'$.
(4) For all $\delta \in \mathbb{R}_+$, there exists $\varepsilon > 0$ such that whenever $\delta' \in [\delta, \delta + \varepsilon]$, we have

$$x \sim_{\delta'} x' \iff x \sim_\delta x', \text{ for all } x, x' \in X.$$

Then, letting $P_X(\delta)$ denote the partition induced by $\sim_\delta$ for each $\delta \in \mathbb{R}_+$, we can verify the following: (1) $P_X(\delta)$ refines $P_X(\delta')$ whenever $\delta' \geqslant \delta \in \mathbb{R}$, (2) $P_X(\delta)$ is the one block partition of X for all $\delta \geqslant \delta_F$, for some $\delta_F \in \mathbb{R}_+$, (3) $P_X(0)$ consists of the singleton partition, and (4) for all $\delta \in \mathbb{R}_+$, there exists $\varepsilon > 0$ such that $P_X(\delta) = P_X(\delta')$ for all $\delta' \in [\delta, \delta + \varepsilon]$.

Before defining the Rips and Dowker nHCMs, we will need a definition. Given any set X and $x, x' \in X$, a *chain* from x to x' is an ordered set $c = \{x_0, \ldots, x_n\}$ of points in X such that $x_0 = x$ and $x_n = x'$. The collection of all such chains is denoted $C_X(x, x')$.

6.1. **The Rips nHCM.** For any network (X, A_X), the Rips nHCM is defined by writing, for each $\delta \in \mathbb{R}_+$,

$$x \sim_\delta^{\mathfrak{R}} x' \iff \text{there exists } c \in C_X(x, x') \text{ such that } [x_i, x_{i+1}] \in \mathfrak{R}_X^\delta \text{ for all } x_i, x_{i+1} \in c, \text{ for all } x, x' \in X.$$

One can then verify that $\sim_\delta^{\mathfrak{R}}$ induces an equivalence relation for each $\delta \in \mathbb{R}_+$, and that the resulting familiy of partitions satisfies the conditions of a dendrogram.

Intuitively, one can interpret a weight $A_X(x, x')$ as the cost of "hopping" from node x to node x' in a network X; with this intuition, two nodes are δ-related in a Rips nHCM (written $x \sim_\delta^{\mathfrak{R}}$) if and only if there is a sequence of nodes starting at x and ending at x' such that the cost of hopping back and forth between consecutive nodes is not greater than δ. We remark that this is precisely the situation for a hierarchical clustering method called the *reciprocal clustering method*, which has already been studied in the existing literature [CMRS13, §5]. Note that in [CMRS13], the reciprocal clustering method has been fully characterized, and has also been shown to be stable. By virtue of the characterization result, one can show that the Rips nHCM is equivalent to the reciprocal clustering method. So the results that hold for reciprocal clustering automatically hold for the Rips nHCM.

6.2. **The Dowker nHCM.** As one would expect, we are also able to define *Dowker source* and *sink* nHCMs. For any network (X, A_X) and any $\delta \in \mathbb{R}_+$, we write:

$$x \sim_\delta^{\mathrm{so}} x' \iff \text{there exists } c \in C_X(x, x') \text{ such that } [x_i, x_{i+1}] \in \mathfrak{D}_{\delta,X}^{\mathrm{so}} \text{ for all } x_i, x_{i+1} \in c, \text{ for all } x, x' \in X,$$

$$x \sim_\delta^{\mathrm{si}} x' \iff \text{there exists } c \in C_X(x, x') \text{ such that } [x_i, x_{i+1}] \in \mathfrak{D}_{\delta,X}^{\mathrm{si}} \text{ for all } x_i, x_{i+1} \in c, \text{ for all } x, x' \in X.$$

Once again, we can verify that the partitions associated to the families $\{\sim_\delta^{\mathrm{so}}\}_{\delta \in \mathbb{R}_+}, \{\sim_\delta^{\mathrm{si}}\}_{\delta \in \mathbb{R}_+}$ satisfy the conditions for a dendrogram.

Recall that one of our main results is Theorem 17, where we show that the persistence diagrams arising from the Dowker source and sink filtrations the same. An analogous but strictly stronger result holds in the case of Dowker sink and source nHCMS, i.e. we are able to show that these two nHCMs are equivalent in the sense that the produce the same dendrograms.

Proposition 26 (Dowker duality for hierarchical clustering)**.** *The Dowker souce and sink hierarchical clustering methods are equivalent. In other words, given any dissimilarity network (X, A_X), any $\delta \in \mathbb{R}_+$, and any two points $x, x' \in X$, we have:*

$$x \sim_\delta^{\mathrm{si}} x' \iff x \sim_\delta^{\mathrm{so}} x'.$$

Proof. Let $x, x' \in X$, and suppose $x \sim_\delta^{\mathrm{si}} x'$. Let $c = \{x = x_0, x_1, \ldots, x_p = x'\}$ be a chain such that $[x_i, x_{i+1}] \in \mathfrak{D}_{\delta,X}^{\mathrm{si}}$ for all $x_i, x_{i+1} \in c$. Then there exists a sequence $c' = \{y_1, y_2, \ldots, y_p\}$ such that:

$$(x_0, y_1), (x_1, y_1), (x_1, y_2), (x_2, y_2), (x_2, y_3), (x_3, y_3), \ldots, (x_{p-1}, y_p), (x_p, y_p) \in R_{\delta,X}.$$

Note that $R_{\delta,X}$ was defined in Equation 2. Now we also have

$$[y_1, y_2], [y_2, y_3], \ldots, [y_{p-1}, y_p] \in \mathfrak{D}_{\delta,X}^{\mathrm{so}},$$

where x_i is a δ-source for $[y_i, y_{i+1}]$, for each $1 \leqslant i \leqslant p-1$. Since $(x_0, x_0), (x_0, y_1) \in R_{\delta,X}$, we also have $[x_0, y_1] \in \mathfrak{D}_{\delta,X}^{\mathrm{so}}$, with x_0 acting as a δ-source. Similarly, we have $(x_p, x_p), (x_p, y_p) \in R_{\delta,X}$, and thus $[y_p, x_p] \in \mathfrak{D}_{\delta,X}^{\mathrm{so}}$, with x_p acting as a δ-source. Hence we have $x \sim_\delta^{\mathrm{so}} x'$, via the chain $(x_0, y_1, y_2, \ldots, y_p, x_p)$. This shows $x \sim_\delta^{\mathrm{so}} x' \implies x \sim_\delta^{\mathrm{si}} x'$. The reverse implication is proved similarly. □

Because of Proposition 26, we may refer to either of the Dowker source and sink nHCMs as the *Dowker nHCM* method. We will denote the corresponding equivalence relations for each $\delta \in \mathbb{R}_+$ by $\sim_\delta^{\mathfrak{D}}$. We emphasize that the result in Proposition 26 cannot be obtained as the 0-dimensional case of Theorem 17, because dendrograms carry strictly more information than 0-dimensional persistence diagrams. By the same token, one can obtain the 0-dimensional case of Theorem 17 as a corollary to Proposition 26.

From the proof of Proposition 26, we can devise the following intuition behind the Dowker nHCM: two nodes x, x' in a Dowker nHCM are δ-related (written $x \sim_\delta^{\mathfrak{D}}$) if and only if there is a sequence of nodes starting at x and ending at x' such that the minimum of the costs of the backward and forward hops between consecutive nodes is not greater than δ. We remark that this is precisely the situation for a hierarchical clustering method known as the *unilateral clustering method* in the existing literature [CMRS13, §10]. The unilateral clustering method was defined, characterized, and also shown to be stable in [CMRS13, Theorem 9, Remark 12]. Using the characterization result of unilateral clustering, one can verify that Dowker nHCM and unilateral clustering are equivalent. By virtue of this equivalence, one can import all of the desirable properties of unilateral clustering into the setting of Dowker nHCM.

To conclude this section, we remark that the notions of connectivity encapsulated in the Dowker nHCM will be revisited in §7.4, where we test our methods on a global migration dataset.

Remark 27. Although we assumed that all the networks in this section were dissimilarity networks, we can also define nHCMs for general networks. The output of an nHCM on a general network will be a *treegram*, which can be thought of as a generalized dendrogram. Details can be found in [SCM16].

7. Implementation and experiments on classification and exploratory data analysis

In this section, we present results of four sets of experiments that we ran using simulated and real-world data. In each experiment, we computed the 0 and 1-dimensional Dowker persistence barcodes of a set of networks. All persistent homology computations were carried out using the `Javaplex` package for Matlab. A full description of Javaplex can be found in [TVJA]. We used $\mathbb{K} = \mathbb{Z}_2$ as the field of coefficients for all our persistence computations. The datasets and software used for our computations are available on `https://research.math.osu.edu/networks/Datasets.html`. A version of our simulated hippocampal networks experiment will appear in [CM16c].

Javaplex comes with a function for returning a representative cycle for each persistence interval. We will use this function heavily in our analysis, especially in §7.2-7.4. One drawback to note is that when there are multiple cycles representing the same persistence interval, we do not have any control over which representative cycle is returned by Javaplex.

All networks in the following experiments were normalized to have weights in the range $[0, 1]$. For each network, we computed Dowker sink complexes at resolutions $\delta = 0.01, 0.02, 0.03, \ldots, 1.00$. This filtration was then passed into Javaplex, which produced 0 and 1-dimensional Dowker persistence barcodes and representative cycles.

7.1. Simulated hippocampal networks.

In the neuroscience literature, it has been shown that as an animal explores a given *environment* or *arena*, specific "place cells" in the hippocampus show increased activity at specific spatial regions, called "place fields" [OD71]. Each place cell shows a *spike* in activity when the animal enters the place field linked to this place cell, accompanied by a drop in activity as the animal moves far away from this place field. To understand how the brain processes this data, a natural question to ask is the following: Is the time series data of the place cell activity, referred to as "spike trains", enough to detect the structure of the arena?

Approaches based on homology [CI08] and persistent homology [DMFC12] have shown positive results in this direction. In [DMFC12], the authors simulated the trajectory of a rat in an arena containing "holes." A simplicial complex was then built as follows: whenever $n+1$ place cells with overlapping place fields fired together, an n-simplex was added. This yield a filtered simplicial complexed indexed by a time parameter. By computing persistence, it was then shown that the number of persistent bars in the 1-dimensional barcode of this filtered simplicial complex would accurately represent the number of holes in the arena.

We repeated this experiment with the following change in methodology: we simulated the movement of an animal, and corresponding hippocampal activity, in arenas with a variety of obstacles. We then induced a directed network from each set of hippocampal activity data, and computed the associated 1-dimensional Dowker persistence diagrams. We were interested in seeing if the bottleneck distances between diagrams arising from similar arenas would differ significantly from the bottleneck distance between diagrams arising from different arenas. To further exemplify our methods, we repeated our analysis after computing the 1-dimensional Rips persistence diagrams from the hippocampal activity networks.

In our experiment, there were five arenas. The first was a square of side length $L = 10$, with four circular "holes" or "forbidden zones" of radius $0.2L$ that the trajectory could not intersect. The other four arenas were those obtained by removing the forbidden zones one at a time. In what follows, we refer to the arenas of each type as *4-hole, 3-hole, 2-hole, 1-hole,* and *0-hole arenas*. For each arena, a random-walk trajectory of 5000 steps was generated, where the animal could move along a square grid with 20 points in each direction. The grid was obtained as a discretization of the box $[0, L] \times [0, L]$, and each step had length $0.05L$. The animal could move in each direction with equal probability. If one or more of these moves took the animal outside the arena (a disallowed move), then the probabilities were redistributed uniformly among the allowed moves. Each trajectory was tested to ensure that it covered the entire arena, excluding the forbidden zones. Formally, we write the time steps as a set $T := \{1, 2, \ldots, 5000\}$, and denote the trajectory as a map $\mathrm{traj} : T \to [0, L]^2$.

For each of the five arenas, 20 trials were conducted, producing a total of 100 trials. For each trial l_k, an integer n_k was chosen uniformly at random from the interval $[150, 200]$. Then n_k place fields of radius $0.05L$ were scattered uniformly at random inside the corresponding arena for each l_k. An illustration of the place field distribution is provided in Figure 12. A spike on a place field was recorded whenever the trajectory would intersect it. So for each $1 \leqslant i \leqslant n_k$, the spiking pattern of cell x_i, corresponding to place field PF_i, was recorded via a function $r_i : T \to \{0, 1\}$ given by:

$$r_i(t) = \begin{cases} 1 & : \mathrm{traj}(t) \text{ intersects } \mathrm{PF}_i, \\ 0 & : \text{otherwise} \end{cases} \qquad t \in T.$$

The matrix corresponding to r_i is called the *raster* of cell x_i. A sample raster is illustrated in Figure 12. For each trial l_k, the corresponding network (X_k, ω_{X_k}) was constructed as follows: X_k consisted of n_k nodes representing place fields, and for each $1 \leqslant i, j \leqslant n_k$, the weight $\omega_{X_k}(x_i, x_j)$ was given by:

$$\omega_{X_k}(x_i, x_j) = 1 - \frac{N_{i,j}(5)}{\sum_{i=1}^{n_k} N_{i,j}(5)},$$
$$\text{where } N_{i,j}(5) = \mathrm{card}\left(\left\{(s,t) \in T^2 : t \in [2, 5000], t - s \in [1, 5], r_j(t) = 1, r_i(s) = 1\right\}\right).$$

In words, $N_{i,j}(5)$ counts the pairs of times $(s, t), s < t$, such that cell x_j spikes at a time t after cell x_i spikes at a time s, and the delay between the two spikes is fewer than 5 time steps. The idea is that if cell x_j frequently fires within a short span of time after cell x_i fires, then place fields PF_i and PF_j are likely to be in close proximity to each other. The column sum of the matrix corresponding to ω_{X_k} is normalized to 1, and so $\omega_{X_k}^\top$ can be interpreted as the transition matrix of a Markov process.

Next, we computed the 1-dimensional Dowker persistence diagrams of each of the 100 networks. Note that $\mathrm{Dgm}_1^{\mathfrak{D}}(\omega_X) = \mathrm{Dgm}_1^{\mathfrak{D}}(\omega_X^\top)$ by Proposition 18, so we are actually obtaining the 1-dimensional Dowker persistence diagrams of transition matrices of Markov processes. We then computed a 100×100 matrix consisting of the bottleneck distances between all the 1-dimensional persistence diagrams. The single linkage dendrogram generated from this bottleneck distance matrix is shown in Figure 13. The labels are in the format `env-<nh>-<nn>`, where `nh` is the number of holes in the arena/environment, and `nn` is the number of place fields. Note that with some exceptions, networks corresponding to the same arena are clustered together. We conclude that the Dowker persistence diagram succeeded in capturing the intrinsic differences between the five classes of networks arising from the five different arenas, even when the networks had different sizes.

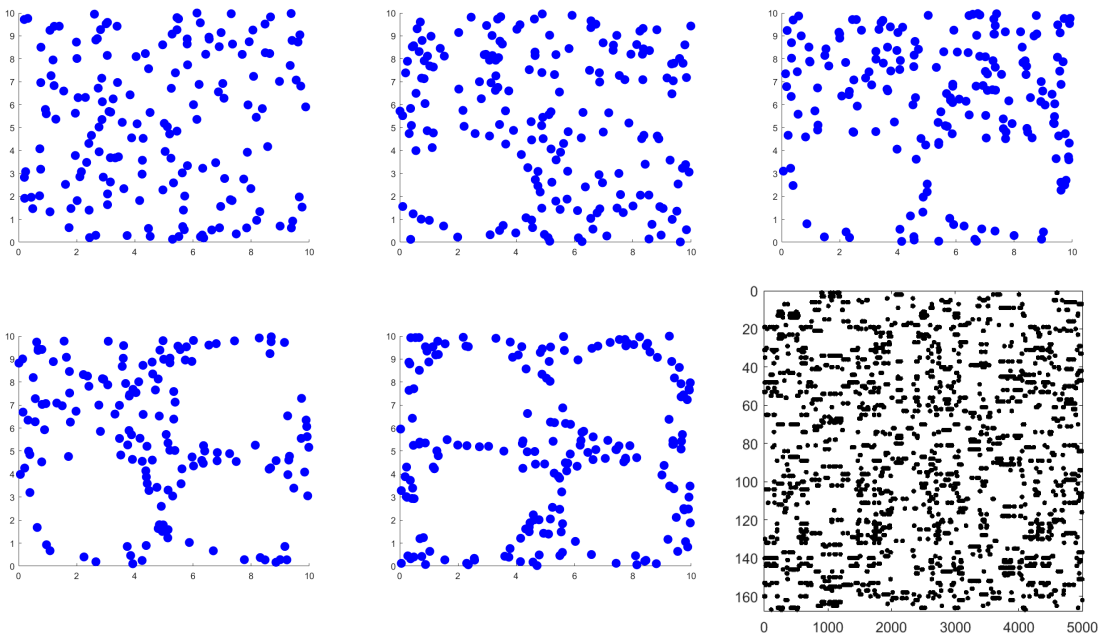

FIGURE 12. **Bottom right:** Sample place cell spiking pattern matrix. The x-axis corresponds to the number of time steps, and the y-axis corresponds to the number of place cells. Black dots represent spikes. **Clockwise from bottom middle:** Sample distribution of place field centers in 4, 3, 0, 1, and 2-hole arenas.

We then computed the Rips persistence diagrams of each network, and computed the 100×100 bottleneck distance matrix associated to the collection of 1-dimensional diagrams. The single linkage dendrogram generated from this matrix is given in Figure 14. Notice that the Rips dendrogram does not do a satisfactory job of classifying arenas correctly.

As a different approach towards comparing the Rips and Dowker methods, we show the 3-dimensional MDS plots of their bottleneck distance matrices in Figure 15. Note that in the Rips MDS plot, all the arena types are scrambled together. In the Dowker MDS plot, the data points corresponding to the 3 and 4-hole arenas are well-separated, and it also appears that the 2-hole arenas are reasonably far apart from the 0 and 1-hole arenas. It seems that the 0 and 1-hole arenas are often confounded with each other, although rarely with the other arenas.

At the moment we are working on expanding the database of hippocampal networks on which to test and evaluate the classifying ability of Dowker persistence diagrams.

Remark 28. We note that an alternative method of comparing the networks obtained from our simulations would have been to compute the pairwise network distances, and plot the results in a dendrogram. But $d_{\mathcal{N}}$ is NP-hard to compute [CM16b], so instead, we are computing the bottleneck distances between 1-dimensional Dowker persistence diagrams, as suggested by Remark 11.

7.2. **U.S. economy input-output accounts.** Each year, the U.S. Department of Commerce Bureau of Economic Analysis (`www.bea.gov`) releases the U.S. input-output accounts, which show the commodity inputs used by each industry to produce its output, and the total commodity output of each industry. Economists use this data to answer two questions: (1) what is the total output of the US economy, and (2) what is the process by which this output is produced and distributed [HP06].

One of the core data types in these accounts is a "make" table, which shows the production of commodities by industries. The industries are labeled according to the North American Industry Classification System (NAICS), and the commodities (i.e. goods or services) produced by each industry are also labeled

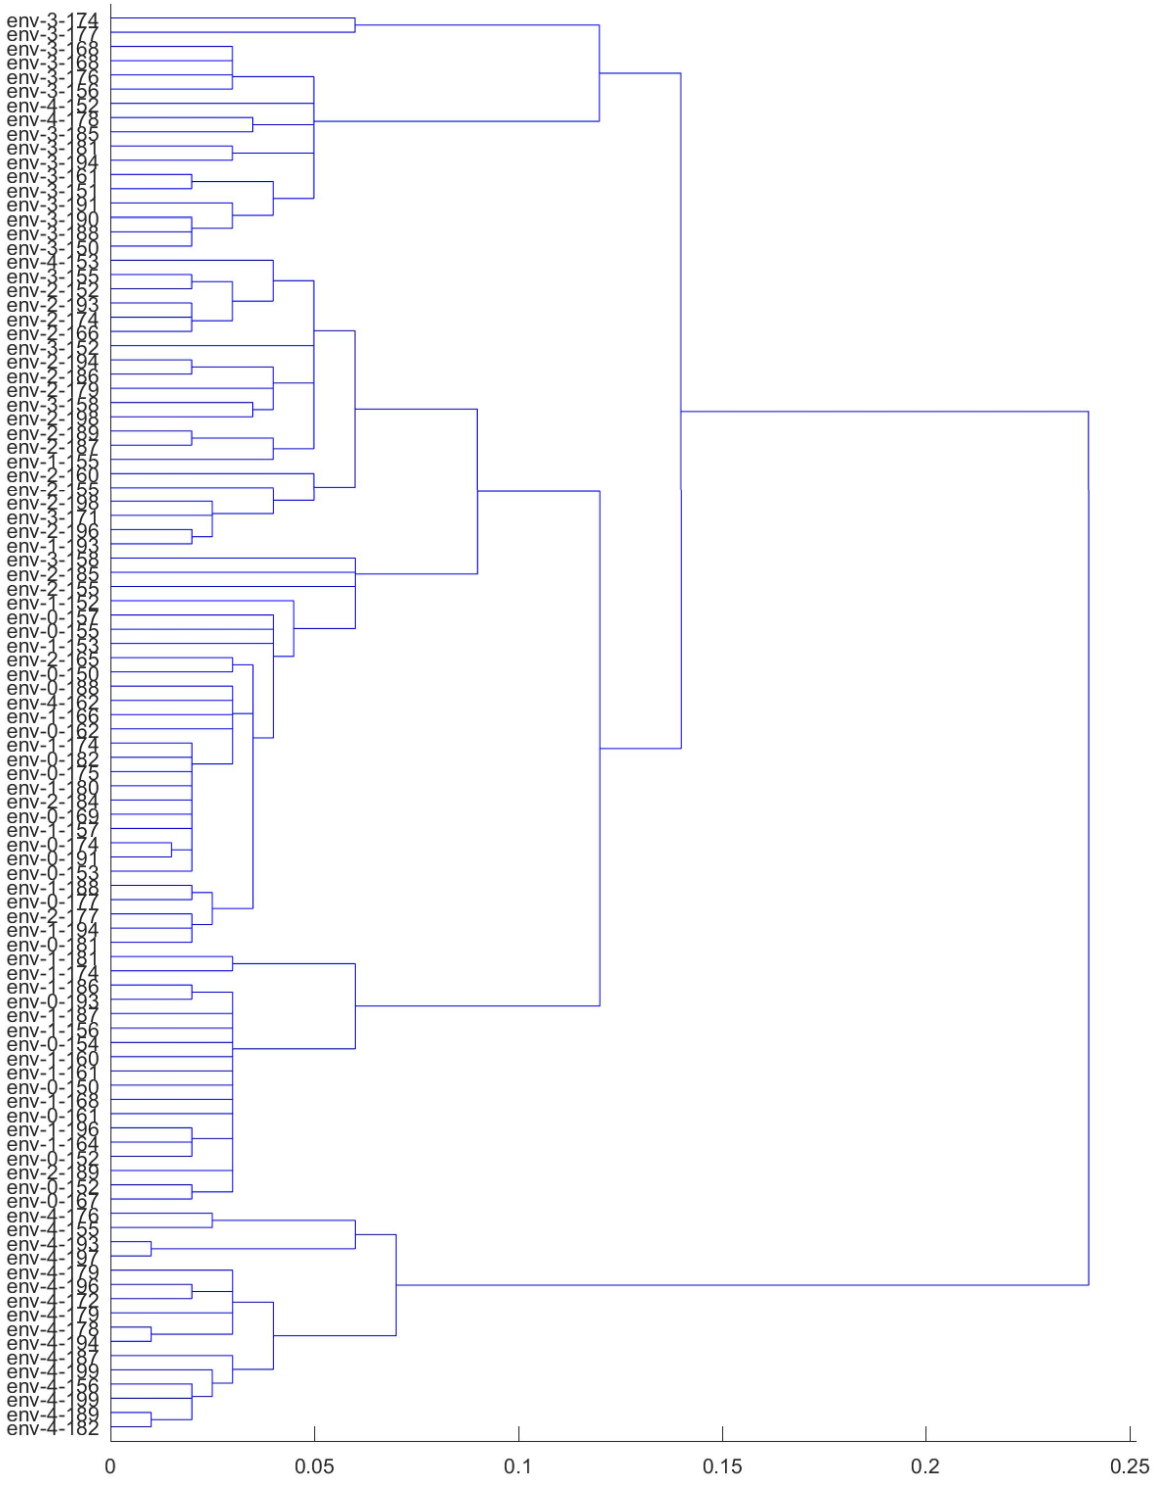

FIGURE 13. Single linkage dendrogram corresponding to the distance matrix obtained by computing bottleneck distances between 1-dimensional Dowker persistence diagrams of our database of hippocampal networks (§7.1). Note that the 4, 3, and 2-hole arenas are well separated into clusters at threshold 0.1.

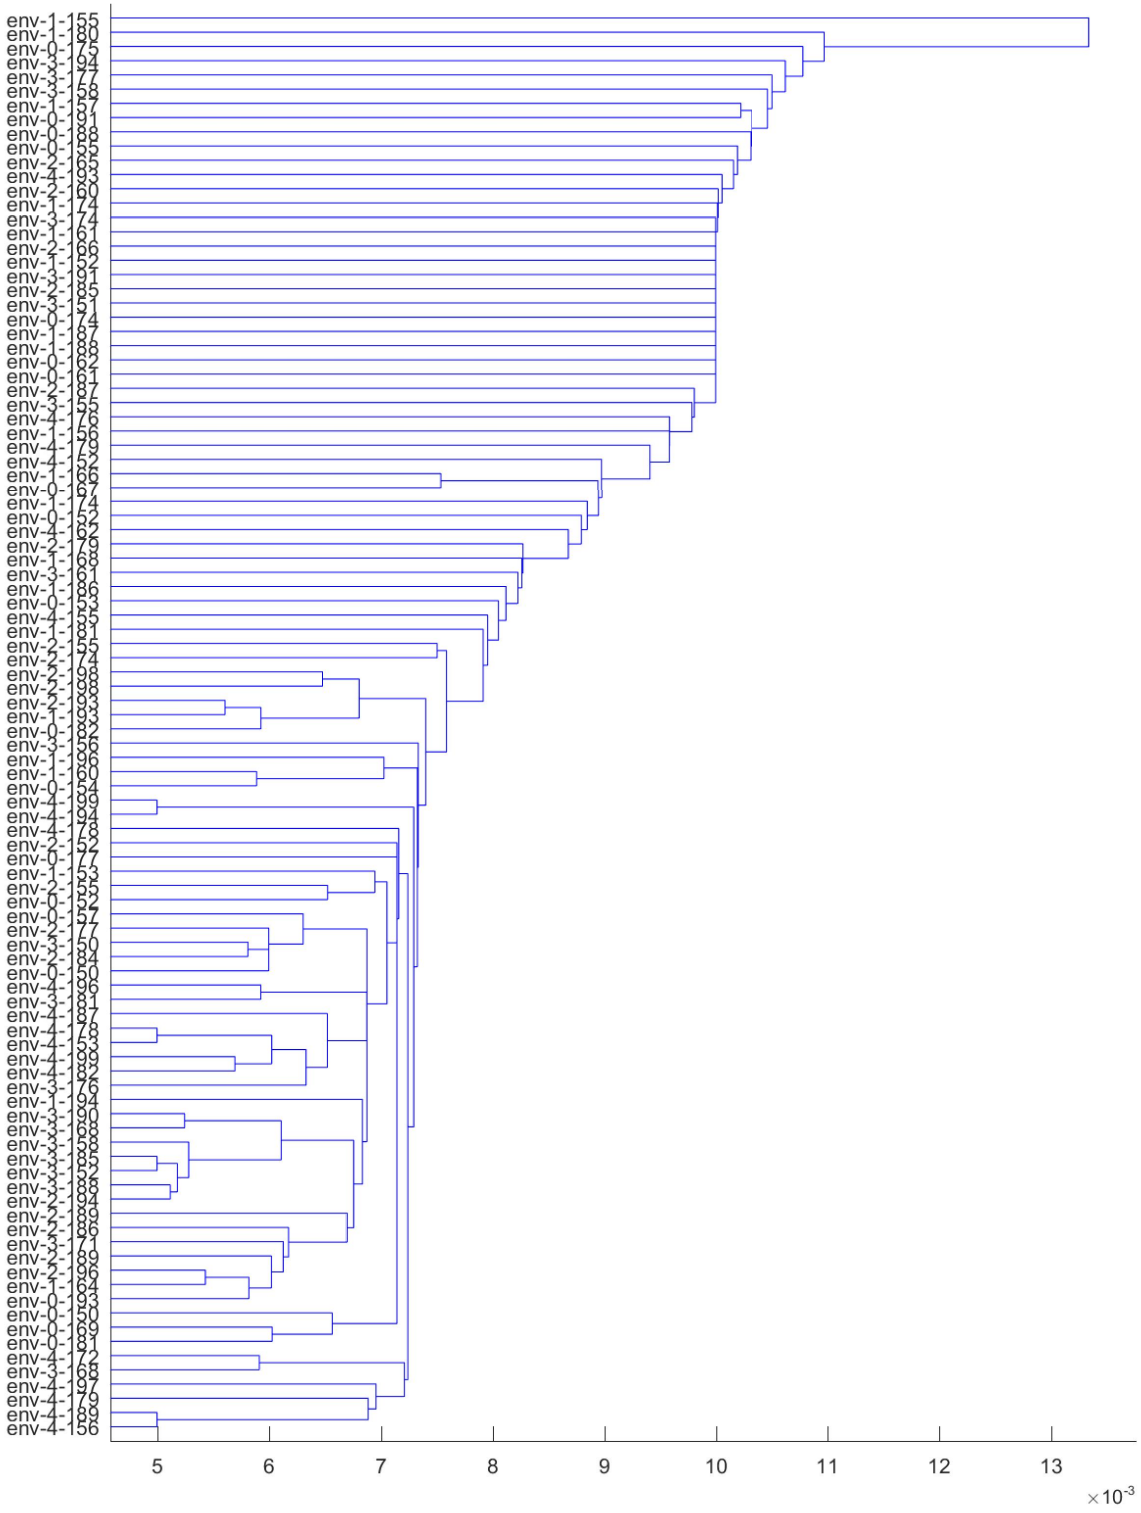

FIGURE 14. Single linkage dendrogram corresponding to the distance matrix obtained by computing bottleneck distances between 1-dimensional Rips persistence diagrams of our database of hippocampal networks (§7.1). Notice that the hierarchical clustering fails to capture the correct arena types.

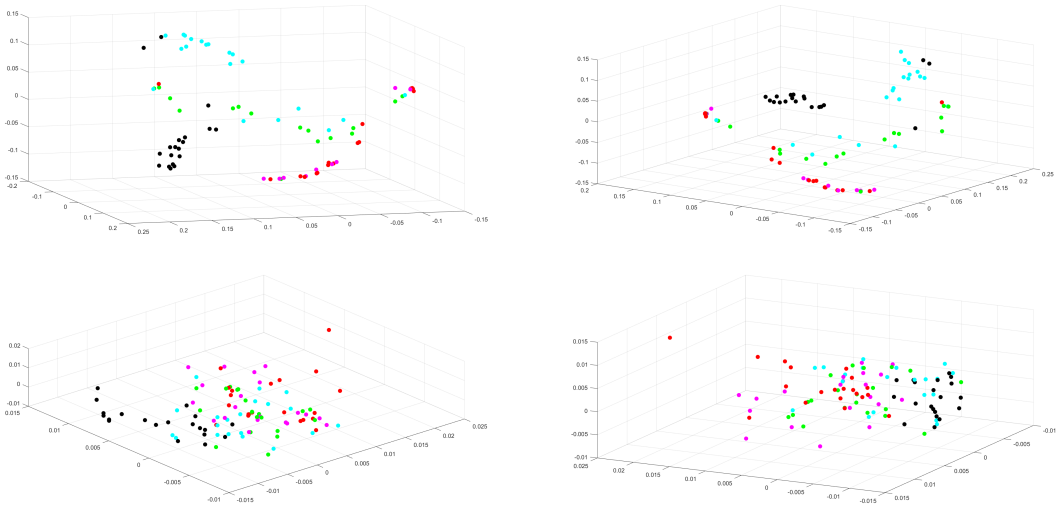

FIGURE 15. MDS scatter plots for our hippocampal networks experiment (§7.1). The magenta, red, green, cyan, and black dots represent 0,1,2,3, and 4-hole arenas, respectively. **Top:** Two different slices of an MDS scatter plot of the 1-dimensional Dowker bottleneck distance matrix. Note that the 3 and 4-hole arenas are well-separated. The 3 and 4-hole arenas appear to be close to each other, but far away from the 0 or 1-hole arenas. The 2-hole arenas are occasionally mixed with the 3-hole arenas, but are typically separated from the 0 and 1-hole arenas. **Bottom:** Two different slices of an MDS scatter plot of the 1-dimensional Rips bottleneck distance matrix. Note that this plot does not appear to contain any discernible structure at all.

according to the NAICS. This make table can be viewed as a network (E, m) consisting of a set of NAICS labels E, and a function $m : E \times E \to \mathbb{Z}_+$. Note that the same labels are used to denote industries and commodities. After fixing an enumeration $(e_1, \ldots, e_n)$ of E, the entry $m(e_i, e_j)$ corresponds to the dollar value (in millions) of commodity type e_j produced by industry e_i. For example, if e corresponds to the economic sector "Farms," then $m(e, e)$ is the dollar value (in millions) of farming commodities produced by the farming industry.

In our next example, we analyze make table data from the U.S. Department of Commerce for the year 2011. Specifically, we begin with a set E of 71 economic sectors, and view it as a network by the process described above. We remark that a complementary data set, the "use" table data for 2011, has been analyzed thoroughly via hierarchical clustering methods in [CMRS13].

Our analysis is motivated by question (2) described above, i.e. what is the process by which commodities are produced and distributed across economic sectors. Note that one can simply read off values from the make table that show *direct* investment between commodities and industries, i.e. which commodities are being produced by which industry. A more interesting question is to find patterns of *indirect* investment, e.g. chains (e_i, e_j, e_k) where industry e_i produces commodities of type e_j, and industry e_j in turn produces commodities of type e_k. Note that if e_i does not produce any commodities of type e_k, then this indirect influence of e_i on e_k is not immediately apparent from the make table. One can manually infer this kind of indirect influence by tracing values across a make table, but this process can become cumbersome when there are large numbers of economic sectors, and when one wants to find chains of greater length.

To automate this process so that finding flows of investment can be used for exploratory data analysis, we take the viewpoint of using persistent homology. Beginning with the 71×71 make table matrix, we first

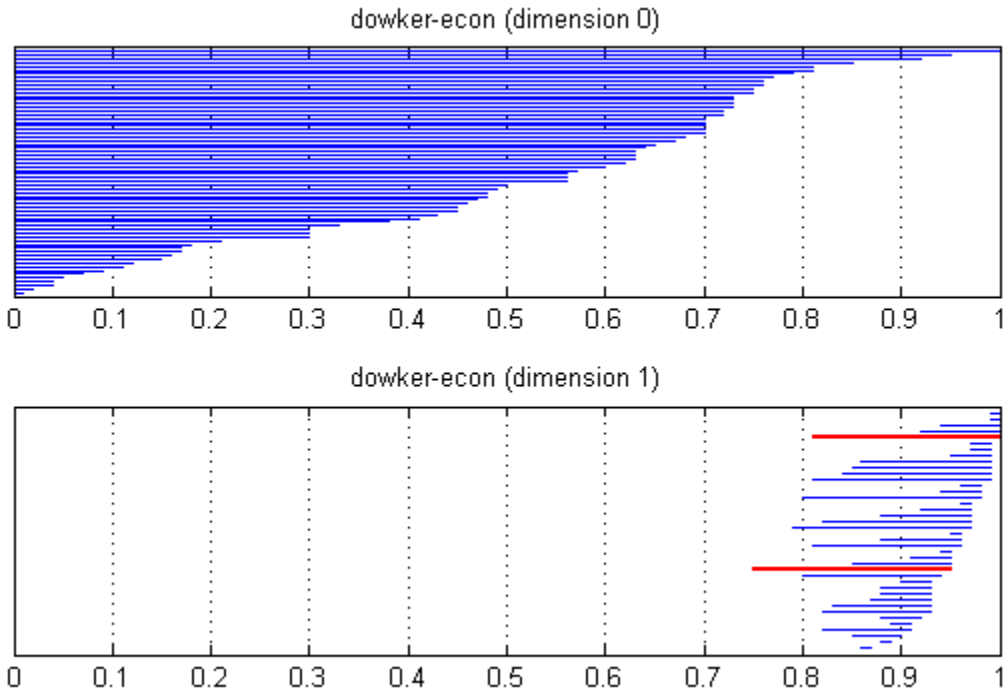

FIGURE 16. 0 and 1-dimensional Dowker persistence barcodes for US economic sector data, obtained by the process described in §7.2. The long 1-dimensional persistence intervals that are colored in red are examined in §7.2.2 and Figures 17,18.

obtained a matrix $\overline{\omega}_E$ defined by:

$$\overline{\omega}_E(e_i, e_j) := \begin{cases} m(e_i, e_j) & : i \neq j \\ 0 & : i = j. \end{cases}, \text{ for each } 1 \leqslant i, j \leqslant 71.$$

Because our goal was to analyze the interdependence of industries and the flow of commodities across industrial sectors, we removed the diagonal as above to discount the commodities produced by each industry in its own type. Next we defined a network (E, ω_E), where ω_E was given by:

$$\omega_E(e_i, e_j) = f\left(\frac{\overline{\omega}_E(e_i, e_j)}{\sum_{e \in E} \overline{\omega}_E(e, e_j)}\right) \text{ for each } 1 \leqslant i, j \leqslant 71.$$

Here $f(x) = 1 - x$ is a function used to convert the original similarity network into a dissimilarity network. The greater the dissimilarity, the weaker the investment, and vice versa. So if $\omega_E(e, e') = 0.85$, then sector e is said to make an investment of 15% on sector e', meaning that 15% of the commodities of type e' produced externally (i.e. by industries other than e') are produced by industry e. After this preprocessing step, we computed the 0 and 1-dimensional Dowker persistence diagrams of the resulting network. The corresponding barcodes are presented in Figure 16, and our interpretation is given below.

7.2.1. *Dependent sectors.* We open our discussion with the 0-dimensional Dowker persistence barcode presented in Figure 16. Recall that Javaplex produces representative 0-cycles for each persistence interval in this 0-dimensional barcode. Typically, these representative 0-cycles are given as the boundary of a 1-simplex, so that we know which pair of sectors merges together into a 1-simplex and converts the 0-cycle into a 0-boundary. We interpret the representative sectors of the shortest 0-dimensional persistence intervals as pairs of sectors where one is strongly dependent on the other. To justify this interpretation, we observe

Sample 0-dimensional persistence invervals from economy data			
Interval $[0, \delta)$	Representative 0-cycle	Labels	δ-sink
$[0.0, 0.1)$	$[47] + [70]$	47=Funds, trusts, and other financial vehicles 70=State and local general government	47
$[0.0, 0.4)$	$[45] + [44]$	45=Securities, commodity contracts, and investments 44=Federal Reserve banks, credit intermediation, and related activities	45
$[0.0, 0.07)$	$[24] + [3]$	3=Oil and gas extraction 24=Petroleum and coal products	24

TABLE 1. The first two columns contain sample 0-dimensional persistence intervals, as produced by Javaplex. We have added the labels in column 3, and the common δ-sinks in column 4.

that the right endpoint of a 0-dimensional persistence interval corresponds to a resolution δ at which two industries e, e' find a common δ-sink. Typically this sink is one of e or e', although this sink is allowed to be a third industry e''. We suggest the following interpretation for being a common δ-sink: of all the commodities of type e'' produced by industries other than e'', over $(1 - \delta) * 100\%$ is produced by each of the industries e, e'. Note that for $\delta < 0.50$, this interpretation suggests that e'' is actually e' (or e), and that over 50% of the commodities of type e' produced by external industries (i.e. by industries other than e') are actually produced by e (resp. e').

In Table 1 we list some sample representative 0-cycles produced by Javaplex. Note that these cycles are representatives for bars that we can actually see on the 0-dimensional barcode in Figure 16. We do not focus any more on finding dependent sectors and direct investment relations, and point the reader to [CMRS13] where this topic has been covered in great detail under the lens of hierarchical clustering (albeit with a slightly different dataset, the "use" table instead of the make table). We also point the reader to §6, where we have provided definitions and interpretations of both the Rips and Dowker network hierarchical clustering methods. In the following subsection, we study some of the persistent 1-dimensional intervals shown in Figure 16, specifically the two longest bars that we have colored in red.

7.2.2. *Patterns of investment.* Examining the representative cycles of the persistent 1-dimensional intervals in Figure 16 allows us to discover patterns of investment that would not otherwise be apparent from the raw data. Javaplex produces representative nodes for each nontrivial persistence interval, so we were able to directly obtain the industrial sectors involved in each cycle. Note that for a persistence interval $[\delta_0, \delta_1)$, Javaplex produces a representative cycle that emerges at resolution δ_0. As more 1-simplices enter the Dowker filtration at greater resolutions, the homology equivalence class of this cycle may coincide with that of a shorter cycle, until finally it becomes the trivial class at δ_1. We have illustrated some of the representative cycles produced by Javaplex in Figures 17 and 18. To facilitate our analysis, we have also added arrows in the figures according to the following rule: for each representative cycle at resolution δ, there is an arrow $e_i \to e_j$ if and only if $\omega_E(e_i, e_j) \leqslant \delta$, i.e. if and only if e_j is a sink for the simplex $[e_i, e_j]$ in $\mathfrak{D}^{\mathrm{si}}_{\delta,E}$.

Consider the 1-dimensional persistence interval $[0.75, 0.95)$, colored in red in Figure 16. The industries involved in a representative cycle for this interval at $\delta = 0.75$ are: Wood products (WO), Primary metals (PM), Fabricated metal products (FM), Petroleum and coal products (PC), Chemical products (CH), and Plastics and rubber products (PL). The entire cycle is illustrated in Figure 17. Starting at the bottom right, note that PC has an arrow going towards CH, suggesting the dependence of the chemical industry on petroleum and coal products. This makes sense because petroleum and coal products are the major organic components used by chemical plants. Chemical products are a necessary ingredient for the synthesis of plastics, which could explain the arrow (CH→PL). Plastic products are commonly used to produce wood-plastic composites, which are low-cost alternatives to products made entirely out of wood. This can explain the arrow PL→WO. Next consider the arrows FM→WO and FM→PM. As a possible interpretation of these arrows, note that fabricated metal frames and other components are frequently used in wood products, and fabricated metal structures are used in the extraction of primary metals from ores. Also note that the

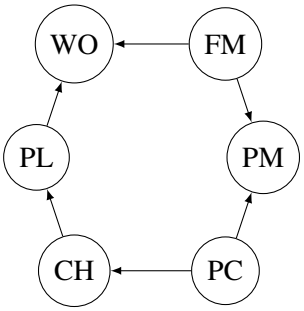

Investment patterns at $\delta = 0.75$

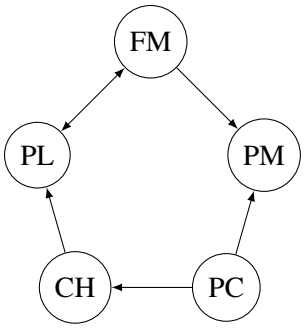

Investment patterns at $\delta = 0.94$

FIGURE 17. Here we illustrate the representative nodes for one of the 1-dimensional persistence intervals in Figure 16. This 1-cycle [PC,CH] + [CH,PL] + [PL,WO] - [WO,FM] + [FM,PM] - [PM,PC] persists on the interval $[0.75, 0.95)$. At $\delta = 0.94$, we observe that this 1-cycle has joined the homology equivalence class of the shorter 1-cycle illustrated on the right. Unidirectional arrows represent an asymmetric flow of investment. A full description of the meaning of each arrow is provided in §7.2.2.

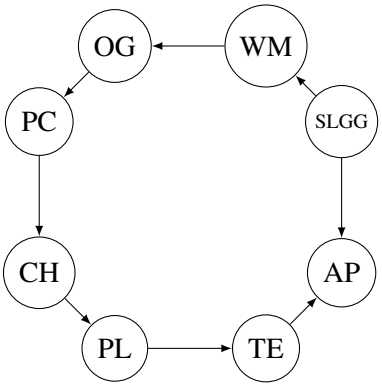

Investment patterns at $\delta = 0.81$

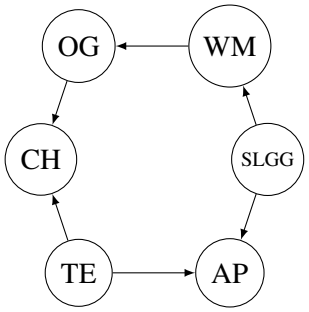

Investment patterns at $\delta = 0.99$

FIGURE 18. Representative nodes for another 1-dimensional persistence interval in Figure 16. A full description of this cycle is provided in §7.2.2.

metal extraction industry is one of the largest consumers of energy. Since energy is mostly produced from petroleum and coal products, this is a possible reason for the arrow PC→PM.

We now consider the 1-dimensional persistence interval $[0.81, 1)$ colored in red in Figure 16. The sectors involved in a representative cycle for this interval at $\delta = 0.81$ are: Petroleum and coal products (PC), Oil and gas (OG), Waste management (WM), State and local general government (SLGG), Apparel and leather and allied products (AP), Textile mills (TE), Plastics and rubber products (PL), and Chemical products (CH). The pattern of investment in this cycle is illustrated in Figure 18, at resolutions $\delta = 0.81$ and $\delta = 0.99$. We have already provided interpretations for the arrows OG→PC→CH→PL above. Consider the arrow PL→TE. This likely reflects the widespread use of polyester and polyester blends in production of fabrics. These fabrics are then cut and sewn to manufacture clothing, hence the arrow TE→AP. Also consider the arrow WM→OG: this suggests the role of waste management services in the oil and gas industry, which makes sense because the waste management industry has a significant role in the treatment and disposal of hazardous materials produced in the oil and gas industry. Finally, note that the arrows SLGG→WM and SLGG→AP likely suggest the dependence of the waste management and apparel industries on state and local government support.

We note that there are numerous other 1-dimensional persistence intervals in Figure 16 that could be worth exploring, especially for economists who regularly analyze the make tables in input-output accounts and are better prepared to interpret this type of data. The results obtained in our analysis above suggest that viewing these tables as asymmetric networks and then computing their persistence diagrams is a reasonable method for uncovering their hidden attributes.

7.3. **U.S. migration.** The U.S. Census Bureau (www.census.gov) publishes data on the annual migration of individuals between the 50 states of the U.S., the District of Columbia, and Puerto Rico. An individual is said to have migrated from State A to State B in year Y if their residence in year Y is in State B, and in the previous year, their residence was in State A. In this section, we define a network structure that encapsulates migration flows within the U.S., and study this network via its Dowker persistence diagrams. We use migration flow data from 2011 in our work. We begin with a set $S = \{s_1, \ldots, s_{52}\}$ of these 52 regions and a function $m : S \times S \to \mathbb{Z}_+$, where $m(s_i, s_j)$ represents the number of migrants moving from s_i to s_j. We define a network (S, ω_S), with ω_S given by:

$$\omega_S(s_i, s_j) = f\left(\frac{m(s_i, s_j)}{\sum_{s_i \in S, i \neq j} m(s_i, s_j)}\right) \text{ if } s_i \neq s_j, \qquad \omega_S(s_i, s_i) = 0, \qquad s_i, s_j \in S,$$

where $f(x) = 1 - x$. The purpose of f is to convert similarity data into dissimilarity data. A large value of $\omega_S(s_i, s_j)$ means that few people move from s_i to s_j. The diagonal is removed to ensure that we focus on migration patterns, not on the base population of each state.

7.3.1. *The Dowker complex of migration.* Because we will eventually compute the Dowker persistence diagram of the network (S, ω_S), we first suggest interpretations of some aspects of the Dowker complex that we construct as an intermediate step. Let $\delta \geqslant 0$, and write $\mathfrak{D}_\delta^{\mathrm{si}} := \mathfrak{D}_{\delta,S}^{\mathrm{si}}, C_2^\delta := C_2(\mathfrak{D}_\delta^{\mathrm{si}})$. We proceed by defining, for any state $s \in S$, the *migrant influx of* s:

$$\mathrm{influx}(s) := \sum_{s' \in S, s' \neq s} m(s', s).$$

Interpretation of a δ-sink. Next we probe the meaning of a δ-sink in the migration context. For simplicity, we first discuss δ-sinks of 1-simplices. Let $[s, s']$ be a 1-simplex for $s, s' \in S$, and let $s'' \in S$ be a δ-sink for $[s, s']$. Then, by unwrapping the definitions above, we see that s'' receives *at least* $(1 - \delta)(\mathrm{influx}(s''))$ migrants from each of s, s'. This suggests the following physical interpretation of the 1-simplex $[s, s']$: in 2010, there were at least $(1 - \delta)(\mathrm{influx}(s''))$ residents in each of s and s' who had a common goal of moving to s'' in 2011. There could be a variety of reasons for interregional migration—people might be moving for employment purposes, for better climate, and so on—but the important point here is that we have a quantitative estimate of residents of s and s' with similar relocation preferences. On the other hand, letting $r, r' \in S$ be states such that $[r, r'] \notin \mathfrak{D}_\delta^{\mathrm{si}}$, the lack of a common δ-sink suggests that residents of r and r' might have significantly different migration preferences. Following this line of thought, we hypothesize the following:

Residents of states that span a 1-simplex in $\mathfrak{D}_\delta^{\mathrm{si}}$ are more similar to each other (in terms of migrational preferences) than residents of states that do not span a 1-simplex.

More generally, when n states form an n-simplex in $\mathfrak{D}_\delta^{\mathrm{si}}$, we say that they exhibit *coherence of preference* at resolution δ. The idea is that the residents of these n states have a mutual preference for a particular attractor state, which acts as a δ-sink. Conversely, a collection of n states that do not form an n-simplex are said to exhibit *incoherence of preference* at resolution δ—residents of these states do not agree strongly enough on a common destination for migration.

Interpretation of a connected component. Now we try to understand the physical interpretation of a *connected component* in $\mathfrak{D}_\delta^{\mathrm{si}}$. Recall that we have already characterized connected components of a Dowker complex in §6; here we provide a treatment specialized to the case of migration. Also recall that two states $s, s' \in S$ belong to a connected component in $\mathfrak{D}_\delta^{\mathrm{si}}$ if and only if there exist $s_1 = s, \ldots, s_n = s' \in S$ such that:

$$[s_1, s_2], [s_2, s_3], [s_3, s_4], \ldots, [s_{n-1}, s_n] \in \mathfrak{D}_\delta^{\mathrm{si}}. \tag{5}$$

Let $s_1, \ldots, s_n \in S$ be such that Condition 5 above is satisfied. Note that this implies that there exists a δ-sink r_i for each $[s_i, s_{i+1}]$, for $1 \leqslant i \leqslant n-1$. Let $r_1, \ldots, r_{n-1}$ be δ-sinks for $[s_1, s_2], \ldots, [s_{n-1}, s_n]$. We can further verify, using the fact that ω_S vanishes on the diagonal, that the sinks $r_1, \ldots, r_{n-1}$ themselves belong to this connected component:

$$[s_1, r_1], [r_1, s_2], [s_2, r_2], [r_2, s_3] \ldots, [r_{n-1}, s_n] \in \mathfrak{D}_\delta^{\mathrm{si}}.$$

The moral of the preceding observations can be summarized as follows:

The vertex set of any connected component of $\mathfrak{D}_\delta^{\mathrm{si}}$ contains a special subset of "attractor" or "sink" states at resolution δ.

So in 2010, for any $i \in \{1, n-1\}$, there were at least $(1-\delta)(\mathrm{influx}(r_i))$ people in s_i and in s_{i-1} who had a common goal of moving to r_i in 2011. Moreover, for any $i, j \in \{1, \ldots, n\}, i \neq j$, there were at least $\min_{1 \leqslant i \leqslant n-1}(1-\delta)(\mathrm{influx}(r_i))$ people in each of s_i and s_j in 2010 who migrated elsewhere in 2011. From a different perspective, we are able to distinguish all the states in a connected component that are significantly attractive to migrants (the sinks/receivers), and we have quantitative estimates on the migrant flow within this connected component into its sink/receiver states.

Consider the special case where each state in a connected component of n states, written as $s_1, s_2, \ldots, s_n$, loses $(1-\delta)(\mathrm{influx}(r))$ residents to a single state $r \in S$. By the preceding observations, r belongs to this connected component, and we can write $r = s_1$ (relabeling as needed). Then we observe that the n states $\{r, s_2, \ldots, s_n\}$ form an n-simplex, with r as a common sink. In this case, we have $\omega_S(s_i, r) \leqslant \delta$ for each $2 \leqslant i \leqslant n$. Also note that if we write

$$\begin{aligned} v_n^\delta &:= [r, s_2] + [s_2, s_3] + [s_3, s_4] + \ldots + [s_{n-1}, s_n] + [s_n, r] \in C_1^\delta \\ \gamma_n^\delta &:= [r, s_2, s_3] + [r, s_3, s_4] + \ldots + [r, s_{n-1}, s_n] \in C_2^\delta, \end{aligned}$$

then we can verify that $\partial_1^\delta(v_n) = 0$, and $\partial_2^\delta(\gamma_n) = v_n$. In other words, we obtain a 1-cycle that is automatically the boundary of a 2-chain, i.e. is trivial upon passing to homology.

In general, a connected component in $\mathfrak{D}_\delta^{\mathrm{si}}$ might contain chains of states that form loops, i.e. states $s_1, s_2, \ldots, s_n$ such that:

$$[s_1, s_2], [s_2, s_3], [s_3, s_4], \ldots, [s_{n-1}, s_n], [s_n, s_1] \in \mathfrak{D}_\delta^{\mathrm{si}}. \tag{6}$$

Note that Condition 6 is of course more stringent than Condition 5. By writing such a loop in the form of v_n^δ above, we can verify that it forms a 1-cycle. Thus a connected component containing a loop will be detected in a 1-dimensional Dowker persistence diagram, unless the resolution at which the 1-cycle appears coincides with that at which it becomes a 1-boundary.

Interpretation of 1-cycles. The preceding discussion shows that it is necessary to determine not just 1-cycles, but also the 1-boundaries that they eventually form. Any 1-boundary arises as the image of ∂_2^δ applied to a linear combination of 2-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$. Note that in this context, each 2-simplex is a triple of states $[s_i, s_j, s_k]$ with a common sink r to which each of s_i, s_j, s_k has lost $(1-\delta)(\mathrm{influx}(r))$ residents between 2010 and 2011. Alternatively, at least $(1-\delta)(\mathrm{influx}(r))$ residents from each of s_i, s_j, s_k had a common preference of moving to r between 2010 and 2011. Next let $\{[s_1, s_1', s_1''], [s_2, s_2', s_2''], \ldots, [s_n, s_n', s_n'']\}$ be a collection of 2-simplices in $\mathfrak{D}_\delta^{\mathrm{si}}$, with sinks $\{r_1, \ldots, r_n\}$. One way to consolidate the information they contain is to simply write them as a sum:

$$\tau_n^\delta := [s_1, s_1', s_1''] + [s_2, s_2', s_2''] + \ldots + [s_n, s_n', s_n''] \in C_2^\delta.$$

Notice that applying the boundary map to τ_n yields:

$$z_n^\delta := \partial_2^\delta(\tau_n) = \sum_{i=1}^{n} \big([s_i', s_i''] - [s_i, s_i''] + [s_i, s_i']\big).$$

At this point we have a list of triples of states, and for each triple we have a quantitative estimate on the number of residents who have a preferred state for migration in common. Now we consider the following relaxation of this situation: for a fixed $i \in \{1, \ldots, n\}$ and some $\delta_0 < \delta$, it might be the case that r_i is no longer a mutual δ_0-sink for $[s_i, s_i', s_i'']$, or even that there is *no* δ_0-sink for $[s_i, s_i', s_i'']$. However, there might still be δ_0-sinks u, u', u'' for $[s_i, s_i'], [s_i', s_i''], [s_i'', s_i]$, respectively. In such a case, we see that $\tau_n^{\delta_0} \notin C_2^{\delta_0}$, but $z_n^{\delta_0} \in C_1^{\delta_0}$. Thus $0 \neq \langle z_n \rangle_{\delta_0} \in H_1(\mathfrak{D}_{\delta_0}^{\mathrm{si}})$. Assuming that $\delta > \delta_0$ is the minimum resolution at which $\langle z_n \rangle_\delta = 0$, we then have a general description of the way in which persistent 1-cycles might arise.

A very special case of the preceding example occurs when we are able to choose a δ-sink r_i for each $[s_i, s_i', s_i'']$, $i \in \{1, \ldots, n\}$, such that $r_1 = r_2 = \ldots = r_n$. In this case, we say that $z_n^{\delta_0}$ becomes a 1-boundary due to a single mutual sink r_1. This situation is illustrated in Figure 19. Also note the interpretation of this special case: assuming that z_n^δ is a 1-boundary, we know that each of the states in the collection $\cup_{i=1}^n \{s_i, s_i', s_i''\}$ loses $(1-\delta)(\mathrm{influx}(r_1))$ residents to r_1 between 2010 and 2011. This signals that r_1 is an especially strong attractor state.

We remark that none of the 1-cycles in the U.S. migration data set that we analyzed exhibited the property of becoming a boundary due to a single mutual sink. However, we did find several examples of this special phenomenon in the global migration dataset studied in §7.4. One of these special sinks turns out to be Djibouti, which is a gateway from the Horn of Africa to the Middle East, and is both a destination and a port of transit for migrants moving between Asia and Africa.

Interpretation of barcodes in the context of migration data. Having suggested interpretations of simplices, cycles, and boundaries, we now turn to the question of interpreting a persistence barcode in the context of migration. Note that when computing persistence barcodes, Javaplex can return a representative cycle for each bar, with the caveat that we do not have any control over which representative is returned. From the 1-dimensional Dowker persistence barcode of a migration dataset, we can use the right endpoint of a bar to obtain a 1-boundary, i.e. a list of triples of states along with quantitative estimates on how many residents from each triple had a preferred migration destination in common. In the special case where the 1-boundary forms due to a single mutual sink, we will have a further quantitative estimate on how many residents from each state in the 1-boundary migrated to the mutual sink. The left endpoint of a bar in the 1-dimensional Dowker persistence barcode corresponds to a special connected component with the structure of a 1-cycle. Notice that all the connected components are listed in the 0-dimensional Dowker persistence diagram. See §7.4 for some additional comments.

Interpretation of error between lower bounds and true migration. In each of Tables 2 and 3 (and Tables 4, 5 in §7.4), we have provided lower bounds on migration flow between certain states, following the discussion above. More precisely, we do the following:

- **0-cycles:** Given a persistence interval $[0, \delta)$, $\delta \in \mathbb{R}$, and a representative 0-cycle, we find the 1-simplex that converts the 0-cycle into a 0-boundary at resolution δ. We then find a δ-sink for this 1-simplex, and estimate a lower bound on the migrant flow into this δ-sink.
- **1-cycles:** Given a persistence interval $[\delta_0, \delta_1)$, $\delta_0, \delta_1 \in \mathbb{R}$, and a representative 1-cycle, we find a δ_0-sink for each 1-simplex in the 1-cycle, and estimate a lower bound on the migrant flow into this δ_0-sink from its associated 1-simplex.

We also provide the true migration flows beside our lower bound estimates. However, in each of our analyses, we incur a certain error between our lower bound and the actual migration value. We now provide some interpretations for this error.

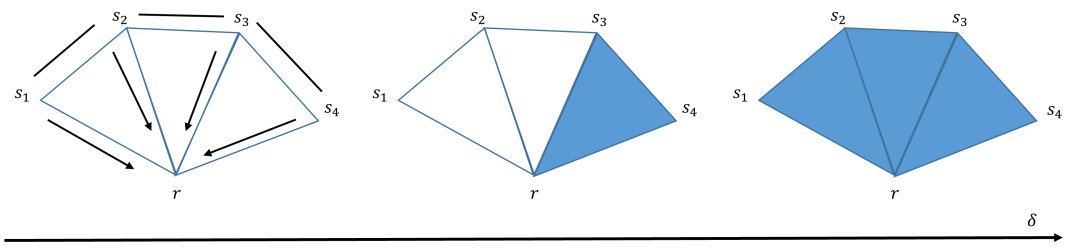

FIGURE 19. An example of a 1-cycle becoming a 1-boundary due to a single mutual sink r, as described in the interpretation of 1-cycles in §7.3.1. The figure on the left shows a connected component of $\mathfrak{D}^{\mathrm{si}}_{\delta,S}$, consisting of $[s_1, s_2], [s_2, s_3], [s_3, s_4]$. The arrows are meant to suggest that r will eventually become a δ-sink for each of these 1-simplices, for some large enough δ. The progression of these simplices for increasing values of δ are shown from left to right. In the leftmost figure, r is not a δ-sink for any of the three 1-simplices. Note that r has become a δ-sink for $[s_3, s_4]$ in the middle figure. Finally, in the rightmost figure, r has become a δ-sink for each of the three 1-simplices.

For the case of 0-cycles, note that all the networks we analyze are normalized to have edge weights in the interval $[0, 1]$. For efficiency, in order to produce a Dowker filtration, we compute $\mathfrak{D}^{\mathrm{si}}_{\delta}$ for δ-values in the set

$$\texttt{delta} := \{0.01, 0.02, 0.03, \ldots, 1.00\}.$$

So whenever we have $\omega_S(s_i, s_j) \notin \texttt{delta}$ for some states $s_i, s_j \in S$, the 1-simplex $[s_i, s_j]$ is not detected until we compute $\mathfrak{D}^{\mathrm{si}}_{\delta'}$, where δ' is the smallest element in $\texttt{delta}$ greater than $\omega_S(s_i, s_j)$. If s_j is a δ-sink in this case, then our predicted lower bound on the migration flow $s_i \to s_j$ will differ by up to $(0.01)(\mathrm{influx}(s_j))$ from the true value. The situation described here best explains the error values in Table 4.

For the case of 1-cycles, we will study a simple motivating example. Suppose we have the following 1-simplices:

$$[s_1, s_2], [s_2, s_3], \ldots, [s_{n-1}, s_n], [s_n, s_1], \text{ for } s_1, \ldots, s_n \in S.$$

For each $i \in \{1, \ldots, n\}$, let $\delta_i \in \mathbb{R}$ denote the resolution at which simplex $[s_i, s_{i+1 \pmod n}]$ emerges. For simplicity, suppose we have $\delta_1 \leqslant \delta_2 \leqslant \delta_3 \leqslant \ldots \leqslant \delta_n$, and also that s_2 is a δ_n-sink for $[s_1, s_2]$. For our lower bound, we estimate that the migrant flow $s_1 \to s_2$ is at least $(1 - \delta_n)(\mathrm{influx}(s_2))$. A better lower bound would be $(1 - \delta_1)(\mathrm{influx}(s_2))$, but the only δ-value that Javaplex gives us access to is δ_n. Because δ_1 could be much smaller than δ_n, it might be the case that our lower bound is much smaller than the true migration.

The preceding discussion suggests the following inference: if a 1-simplex $[s_i, s_{i+1}]$ exhibits a large error between the true migration into a δ_n-sink and the predicted lower bound, then $[s_i, s_{i+1}]$ likely emerged at a resolution proportionately smaller than δ_n. Thus we can interpret the states s_i, s_{i+1} as exhibiting relatively strong coherence of preference. Conversely, 1-simplices that exhibit a smaller error likely emerged at a resolution closer to δ_n—the states forming such 1-simplices exhibited incoherence of preference for a greater range of resolutions. Note that even though we made some simplifying assumptions in our choice of a 1-cycle, a similar analysis can be done for any 1-cycle.

7.3.2. *Analysis of U.S. migration.* We computed the 0 and 1-dimensional Dowker persistence barcodes of the network (S, ω_S) obtained in §7.3. The result is shown in Figure 20. As particular cases, we study the two 1-dimensional bars that are highlighted in red in Figure 20. We obtain representative cycles for each of these bars, and superimpose them on a map of the U.S. in Figure 21. In the remainder of this section, we will discuss the cycles presented in Figure 21.

Analysis of OH-KY-GA-FL cycle			
1-simplex	0.90-sinks	Estimated lower bound on migration	True migration
[FL,OH]	WV	$(1-0.90)(\text{influx}(\text{WV})) = 4597$	$m(\text{FL},\text{WV}) = 4964$ $m(\text{OH},\text{WV}) = 7548$
[FL,GA]	AL	$(1-0.90)(\text{influx}(\text{AL})) = 10684$	$m(\text{FL},\text{AL}) = 12635$ $m(\text{GA},\text{AL}) = 18799$
	GA	$(1-0.90)(\text{influx}(\text{GA})) = 24913$	$m(\text{FL},\text{GA}) = 38658$
[KY,OH]	KY	$(1-0.90)(\text{influx}(\text{KY})) = 9925$	$m(\text{OH},\text{KY}) = 12744$
[GA,KY]	TN	$(1-0.90)(\text{influx}(\text{TN})) = 15446$	$m(\text{GA},\text{TN}) = 16898$ $m(\text{KY},\text{TN}) = 16852$
2-simplex	0.94-sinks	Estimated lower bound on migration	True migration
[FL,OH,KY]	IN	$(1-0.94)(\text{influx}(\text{IN})) = 8640$	$m(\text{FL},\text{IN}) = 11472$ $m(\text{OH},\text{IN}) = 11588$ $m(\text{KY},\text{IN}) = 11071$
	OH	$(1-0.94)(\text{influx}(\text{OH})) = 12363$	$m(\text{FL},\text{OH}) = 18191$ $m(\text{KY},\text{OH}) = 19617$
[FL,GA,KY]	TN	$(1-0.94)(\text{influx}(\text{TN})) = 9268$	$m(\text{FL},\text{TN}) = 10451$ $m(\text{GA},\text{TN}) = 16898$ $m(\text{KY},\text{TN}) = 16852$

TABLE 2. Quantitative estimates on migrant flow, following the interpretation presented in §7.3.1. In each row, we list a simplex of the form $[s_i, s_j]$ (resp. $[s_i, s_j, s_l]$ for 2-simplices) and any possible δ-sinks s_k. We hypothesize that s_k receives at least $(1-\delta)(\text{influx}(s_k))$ migrants from each of s_i, s_j (resp. s_i, s_j, s_l)—these lower bounds are presented in the third column. The fourth column contains the true migration numbers. Notice that the [FL,GA] simplex appears to show the greatest error between the lower bound and the true migration. Following the interpretation suggested earlier in §7.3.1, this indicates that Florida and Georgia appear to have strong coherence of preference, relative to the other pairs of states spanning 1-simplices in this table.

The (OH-KY-GA-FL) cycle. As a representative 1-cycle for the 1-dimensional interval $[0.90, 0.94)$, Javaplex returns the 1-cycle ([FL,OH] + [FL,GA] + [KY,OH] + [GA,FL]). This Ohio-Kentucky-Georgia-Florida cycle first emerges in $C_1^{0.90}$, and becomes a 1-boundary in $C_1^{0.94}$. In Table 2, we provide the 0.90-sinks for each of the 1-simplices, the 0.94-sinks for some 2-simplices that convert this cycle into a boundary, our estimates on the number of migrants to each sink, and also the true numbers of migrants.

The (WA-OR-CA-AZ-UT-ID) cycle. As a representative 1-cycle for the 1-dimensional interval $[0.87, 0.92)$, Javaplex returns the 1-cycle ([CA,OR] + [OR,WA] + [AZ,UT] + [ID,UT] + [ID,WA] + [AZ,CA]). This Washington-Oregon-California-Arizona-Utah-Idaho cycle first emerges in $C_1^{0.87}$, and becomes a 1-boundary in $C_1^{0.92}$. In Table 3, we provide the 0.87-sinks for each of the 1-simplices, the 0.92-sinks that convert this 1-cycle into a boundary, our estimates on the number of migrants to each sink, and the true migration numbers.

To probe the sociological aspects of a 1-cycle, we recall our hypothesis that residents of states that are not connected by a 1-simplex are less similar to each other than residents of states that are connected as such. The West Coast cycle given above seems to follow this hypothesis: It seems reasonable to think that residents of California would be quite different from residents of Idaho or Utah, and possibly quite similar to those of Oregon. Similarly, one would expect a large group of people from Ohio and Kentucky to be quite similar, especially with Cincinnati being adjacent to the state border with Kentucky. The Ohio-Florida simplex might be harder to justify, but given the very small population of their mutual sink West Virginia, it might be the case that the similarity between Ohio and Florida is being overrepresented.

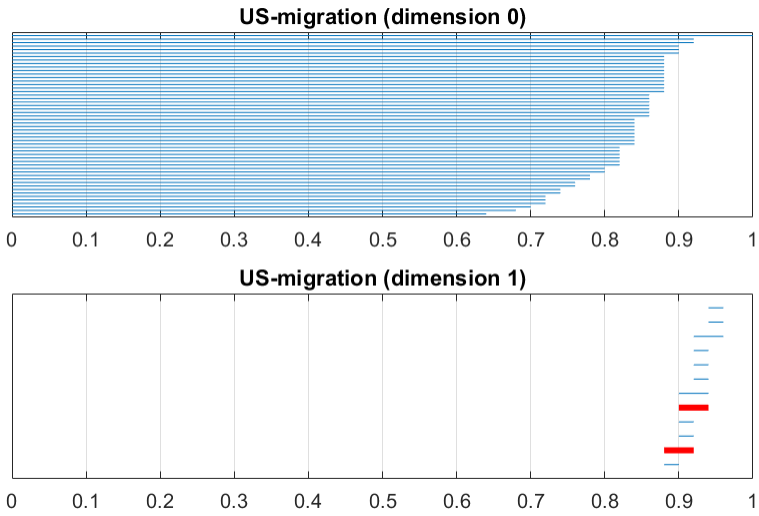

FIGURE 20. 0 and 1-dimensional Dowker persistence barcodes for U.S. migration data

FIGURE 21. U.S. map with representative cycles of the persistence intervals that were highlighted in Figure 20. The cycle on the left appears at $\delta_1 = 0.87$, and the cycle on the right appears at $\delta_2 = 0.90$. The red lines indicate the 1-simplices that participate in each cycle. Each red line is decorated with an arrowhead $s_i \to s_j$ if and only if s_j is a sink for the simplex $[s_i, s_j]$. The blue arrows point towards all possible alternative δ-sinks, and are interpreted as follows: Tennessee is a 0.90-sink for the Kentucky-Georgia simplex, West Virginia is a 0.90-sink for the Ohio-Florida simplex, and Alabama is a 0.90-sink for the Georgia-Florida simplex.

Analysis of WA-OR-CA-AZ-UT-ID cycle			
1-simplex	0.87-sinks	Estimated lower bound on migration	True migration
[CA,OR]	OR	$(1-0.87)(\text{influx}(\text{OR})) = 14273$	$m(\text{CA},\text{OR}) = 18165$
[OR,WA]	OR	$(1-0.87)(\text{influx}(\text{OR})) = 14273$	$m(\text{WA},\text{OR}) = 29168$
[AZ,UT]	UT	$(1-0.87)(\text{influx}(\text{UT})) = 9517$	$m(\text{AZ},\text{UT}) = 10577$
[ID,UT]	ID	$(1-0.87)(\text{influx}(\text{ID})) = 7519$	$m(\text{UT},\text{ID}) = 7538$
[ID,WA]	ID	$(1-0.87)(\text{influx}(\text{ID})) = 7519$	$m(\text{WA},\text{ID}) = 10895$
[AZ,CA]	AZ	$(1-0.87)(\text{influx}(\text{AZ})) = 27566$	$m(\text{CA},\text{AZ}) = 35650$
2-simplex	0.92-sinks	Estimated lower bound on migration	True migration
[OR,WA,UT]	ID	$(1-0.92)(\text{influx}(\text{ID})) = 4627$	$m(\text{OR},\text{ID}) = 6236$ $m(\text{WA},\text{ID}) = 10895$ $m(\text{UT},\text{ID}) = 7538$
[AZ,CA,ID]	UT	$(1-0.92)(\text{influx}(\text{UT})) = 5856$	$m(\text{AZ},\text{UT}) = 10577$ $m(\text{CA},\text{UT}) = 8944$ $m(\text{ID},\text{UT}) = 6059$

TABLE 3. Quantitative estimates on migrant flow, following the interpretation presented in §7.3.1. The entries in this table follow the same rules as those of Table 2. Notice that the [OR,WA] and [AZ,CA] simplices show the greatest error between the lower bound and the true migration. Following the interpretation in §7.3.1, this suggests that these two pairs of states exhibit stronger coherence of preference than the other pairs of states forming 1-simplices in this table.

Our analysis shows that by using Dowker persistence diagrams for exploratory analysis of migration data, we can obtain meaningful lower bounds on the number of residents from different states who share a common migration destination. An interesting extension of this experiment would be to study the persistence barcodes of migration networks over a longer range of years than the 2010-2011 range that we have used here: ideally, we would be able to detect changing trends in migration from changes in the lower bounds that we obtain.

7.4. **Global migration.** For our next example, we study data from the World Bank Global Bilateral Migration Database [ÖPSW11] on global bilateral migration between 1990 and 2000. This dataset is available on `http://databank.worldbank.org/`. As in the case of the U.S. migration database, we begin with a set $C = \{c_1, c_2, \ldots, c_{231}\}$ of 231 global regions and a function $m : C \times C \to \mathbb{Z}_+$, where $m(c_i, c_j)$ represents the number of migrants moving from region c_i to region c_j. We define a network (C, ω_C), with ω_C given by:

$$\omega_C(c_i, c_j) = f\left(\frac{m(c_i, c_j)}{\sum_{c_i \in C, i \neq j} m(c_i, c_j)}\right) \text{ if } c_i \neq c_j, \qquad \omega_C(c_i, c_i) = 0, \qquad c_i, c_j \in C,$$

where $f(x) = 1 - x$. The 0 and 1-dimensional Dowker persistence barcodes that we obtain from this network are provided in Figure 22. Some of the 0 and 1-dimensional persistence intervals are tabulated in Tables 4 and 5.

We interpret connected components, simplices, cycles and boundaries for the Dowker sink complexes constructed from the global migration data just as we did for the U.S. migration data in §7.3.1.

We draw the reader's attention to two interesting features in the persistence barcodes of the global migration dataset. First, the 0-dimensional barcode contains many short bars (e.g. bars of length less than 0.2). In contrast, the shortest bars in the 0-dimensional barcode for the U.S. migration data had length greater than 0.65. In our interpretation, which we explain more carefully below, this observation suggests that migration patterns in the U.S. are relatively uniform, whereas global migration patterns can be more skewed. Second,

because there are many more 1-dimensional persistence intervals, it is easier to find a 1-cycle that becomes a boundary due to a single mutual sink, i.e. due to an especially strong "attractor" region.

For the first observation, consider a 0-dimensional persistence interval $[0, \delta)$, where δ is assumed to be small. Formally, this interval represents the persistence of a 0-cycle that emerges at resolution 0, and becomes a 0-boundary at resolution δ. One can further verify the following: this interval represents the resolutions for which a 0-simplex $[c_i], c_i \in C$ remains disconnected from other 0-simplices, and δ is a resolution at which c_i forms a 1-simplex with some $c_j \in C$. Recall from §7.3.1 that this means the following: either there exists a region $c_k \notin \{c_i, c_j\}$ which receives at least $(1-\delta)(\text{influx}(c_k))$ migrants from each of c_i and c_j, or $c_k = c_i$ (or c_j) and c_k receives over $(1-\delta)(\text{influx}(c_k))$ migrants from c_j (resp. c_i). The first case cannot happen when $\delta < 0.5$, because this would mean that c_k receives strictly over 50% of its migrant influx from each of c_i and c_j. Thus when $\delta < 0.5$, we know that $c_k = c_i$ (or $c_k = c_j$), and c_k receives over 50% of its migrant influx from c_j (resp. c_i). For very small δ, we then know that *most* of the migrants into c_k arrive from c_j (resp. c_i).

For convenience, let us assume that $\delta < 0.2$, that $c_k = c_i$, and that c_k receives over 80% of its migrant influx from c_j. This might occur for a variety of reasons, some of which are: (1) there might be war or political strife in c_j and c_k might be letting in refugees, (2) c_k might have achieved independence or administrative autonomy and some residents from c_j might be flocking to c_k because they perceive it to be their homeland, and (3) c_j might be overwhelmingly populous in comparison to other neighboring regions of c_k, so that the contribution of c_j to the migrant influx of c_k dominates that of other regions.

Notice that neither of the first two reasons listed above are valid in the case of U.S. migration. The third reason is valid in the case of a few states, but nevertheless, the shortest 0-dimensional persistence interval in the U.S. migration dataset has length greater than 0.65. In other words, the minimal resolution at which a 1-simplex forms in the U.S. migration data is 0.65. This in turn means that there is no state in the U.S. which receives over 35% of its migrant influx from any single other state. Based on this reasoning, we interpret the migration pattern of the U.S. as "diffuse" or "uniform", and that of the world as a whole as "skewed" or "biased". This makes intuitive sense, because despite the heterogeneity of the U.S. and differences in state laws and demographics, any resident can easily migrate to any other state of their choice while maintaining similar legal rights, salary, and living standards.

In Table 4, we list some short 0-dimensional persistence intervals for the global migration dataset. For each interval $[0, \delta)$, we also include the 1-simplex that emerges at δ, the δ-sink associated to this 1-simplex, and our lower bound on the migrant influx into this sink. Note that the error between the true migration numbers and our predicted lower bounds is explained in §7.3.1. Also notice that many of the migration patterns provided in Table 4 seem to fit with the suggestions we made earlier: (1) political turmoil in the West Bank and Gaza (especially following the Gulf War) prompted many Palestinians to enter Syria, (2) Greenland and Macao are both autonomous regions of Denmark and China, respectively, and (3) India's population far outstrips that of its neighbors, and its migrant flow plays a dominating role in the migrant influx of its neighboring states.

For the second observation, recall from our discussion in §7.3.1 that whenever we have a 1-cycle involving regions $c_1, \ldots, c_n$ that becomes a 1-boundary at resolution $\delta \geqslant 0$ due to a single mutual sink c_{n+1}, we know that c_{n+1} receives at least $(1-\delta)(\text{influx}(c_{n+1}))$ migrants from each of $c_1, \ldots, c_n$. As such, c_{n+1} can be perceived to be an especially strong attractor region. In Table 5 we list some 1-cycles persisting on an interval $[\delta_0, \delta_1)$, their mutual δ_1-sinks, our lower bound on migration flow, and the true migration numbers. The reader is again encouraged to check that the true migration agrees with the lower bounds that we predicted. We remark that the first row of this table contains a notable example of a strong attractor region: Djibouti. Djibouti is geographically located at a crossroads of Asia and Africa, and is a major commercial hub due to its access to the Red Sea and the Indian Ocean. As such, one would expect it to be a destination for many migrants in the Horn of Africa, as well as a transit point for migrants moving between Africa and the Middle East.

The Oceania cycle listed in the fourth row of Table 5 can likely be discarded; the very small migrant influx of Samoa indicates that it attractiveness as a sink state is being overrepresented. The second row lists

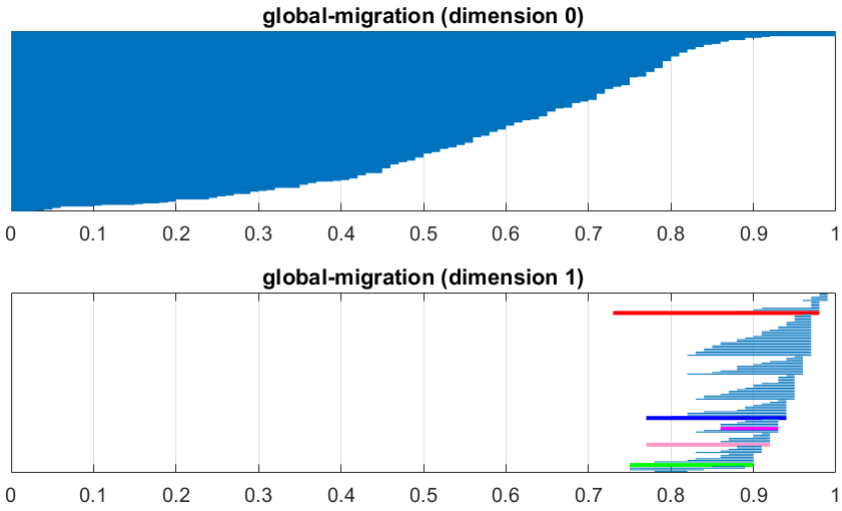

FIGURE 22. Dowker persistence barcodes for global migration dataset.

Short bars in 0-dimensional global migration barcode				
Interval $[0, \delta)$	1-simplex	δ-sink(s)	Lower bound on migration	True migration
[0.0,0.03)	[India, Sri Lanka]	Sri Lanka	$(0.97)(\text{influx}(\text{LKA})) = 383034$	$m(\text{IND}, \text{LKA}) = 384789$
[0.0,0.04)	[India, Bangladesh]	Bangladesh	$(0.96)(\text{influx}(\text{BGD})) = 927270$	$m(\text{IND}, \text{BGD}) = 936151$
[0.0,0.05)	[India, Nepal]	Nepal	$(0.95)(\text{influx}(\text{NPL})) = 927270$	$m(\text{IND}, \text{BGD}) = 936151$
[0.0,0.05)	[India, Pakistan]	Pakistan	$(0.95)(\text{influx}(\text{PAK})) = 2508882$	$m(\text{IND}, \text{PAK}) = 2512906$
[0.0,0.06)	[India, Bhutan]	Bhutan	$(0.94)(\text{influx}(\text{BTN})) = 30206$	$m(\text{IND}, \text{BTN}) = 30431$
[0.0,0.06)	[Denmark, Greenland]	Greenland	$(0.94)(\text{influx}(\text{GRL})) = 6792$	$m(\text{DNK}, \text{GRL}) = 6808$
[0.0,0.10)	[Greece, Albania]	Albania	$(0.90)(\text{influx}(\text{ALB})) = 67008$	$m(\text{GRC}, \text{ALB}) = 67508$
[0.0,0.11)	[Timor-Leste, Indonesia]	Timor-Leste	$(0.89)(\text{influx}(\text{TLS})) = 8246$	$m(\text{IDN}, \text{TLS}) = 8334$
[0.0,0.15)	[West Bank and Gaza, Syria]	Syria	$(0.85)(\text{influx}(\text{SYR})) = 455515$	$m(\text{PSE}, \text{SYR}) = 458611$
[0.0,0.16)	[Macao, China]	Macao	$(0.84)(\text{influx}(\text{MAC})) = 201840$	$m(\text{CHN}, \text{MAC}) = 203877$

TABLE 4. Short 0-dimensional Dowker persistence intervals capture regions which receive most of their incoming migrants from a single source. Each interval $[0, \delta)$ corresponds to a 0-simplex which becomes subsumed into a 1-simplex at resolution δ. We list these 1-simplices in the second column, and their δ sinks in the third column. The definition of a δ-sink enables us to produce a lower bound on the migration into each sink, which we provide in the fourth column. We also list the true migration numbers in the fifth column, and the reader can consult §7.3.1 for our explanation of the error between the true migration and the lower bounds on migration.

China as a strong attractor, which is reasonable given its economic growth between 1990 and 2000, and as a consequence, its attractiveness to foreign workers from neighboring countries. The third row lists Vietnam as a strong sink, and one reason could be that in the 1990s, many refuges who had been displaced due to the Vietnam War were returning to their homeland.

We also illustrate the emergence at δ_0 for some of these cycles in Figure 23.

8. DISCUSSION

In this paper, we have provided a complete description of the Rips and Dowker persistence diagrams of general networks. The stability results we have provided give quantitative guarantees on the robustness of these persistence diagrams. We have provided examples suggesting that Dowker persistence diagrams are

1-cycles with single mutual sink in global migration data				
Interval $[\delta_0, \delta_1)$	Regions involved	Mutual δ_1-sink(s)	Lower bound on migration	True migration
[0.73,0.98)	Djibouti—Ethiopia—Eritrea—Uganda—Somalia	Djibouti	$(0.02)(\text{influx}(\text{DJI})) = 1738$	$m(\text{ERI}, \text{DJI}) = 3259$ $m(\text{ETH}, \text{DJI}) = 25437$ $m(\text{SOM}, \text{DJI}) = 41968$ $m(\text{UGA}, \text{DJI}) = 1811$
[0.77,0.94)	China—Thailand—Philippines	China	$(0.06)(\text{influx}(\text{CHN})) = 12858$	$m(\text{THA}, \text{CHN}) = 14829$ $m(\text{PHL}, \text{CHN}) = 17828$
[0.75,0.89)	China—Indonesia—Malaysia	Vietnam	$(0.11)(\text{influx}(\text{VNM})) = 4465$	$m(\text{CHN}, \text{VNM}) = 8940$ $m(\text{IDN}, \text{VNM}) = 10529$ $m(\text{MYS}, \text{VNM}) = 4813$
[0.86,0.93)	American Samoa—New Zealand—Samoa—Australia	Samoa	$(0.07)(\text{influx}(\text{WSM})) = 397$	$m(\text{ASM}, \text{WSM}) = 1920$ $m(\text{NZL}, \text{WSM}) = 1803$ $m(\text{AUS}, \text{WSM}) = 404$
[0.77,0.92)	Kiribati—Papua New Guinea—Australia—United Kingdom—Tuvalu	Not applicable	Not Applicable	Not Applicable

TABLE 5. Representative 1-cycles for several intervals in the 1-dimensional persistence barcode for the global migration dataset. Each of the first four cycles has the special property that it becomes a boundary due to a single sink at the right endpoint of its associated persistence interval. This permits us to obtain a lower bound on the migration into this sink from each of the regions in the cycle. The last row contains a cycle without this special property. The font colors of the persistence intervals correspond to the colors of the highlighted 1-dimensional bars in Figure 22.

an appropriate method for analyzing general asymmetric networks. For a particular class of such examples, the family of cycle networks, we have fully characterized the 1-dimensional Dowker persistence diagrams. We have provided a proof of Dowker duality, as well as an analogous but strictly stronger result for Dowker hierarchical clustering methods. Finally, we have implemented our methods on a variety of simulated and real-world datasets, and provided interpretations for our results.

We believe that the story of "persistent homology of asymmetric networks" has more aspects to be uncovered. Of particular interest to us is the analysis of alternative methods of producing simplicial complexes from asymmetric networks, for example, the *directed flag complex* construction of [DHR+16]. Some other interesting questions relate to cycle networks: we would like to obtain (1) a characterization of the k-dimensional Dowker persistence diagrams of cycle networks for $k \geqslant 2$, and (2) a characterization of the Rips persistence diagrams of cycle networks for any dimension $k \geqslant 1$.

We perceive that combining the results in our work with the established literature on modeling social phenomena via topological methods [Atk75] can be a powerful method for obtaining new insights in the social sciences. It is our belief that just as persistent homology has made a marked societal and academic impact via its application to metric data, it will continue to do so in the context of network analysis.

Acknowledgments. This work was supported by NSF grants IIS-1422400 and CCF-1526513. We thank Santiago Segarra for providing the U.S. economy dataset used in §7.2. We also thank Pascal Wild for correcting an error on an earlier version of this paper.

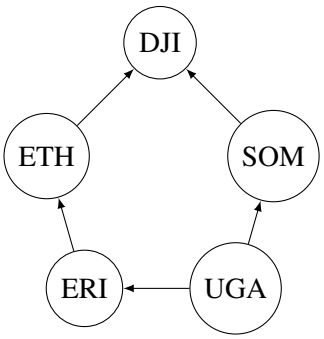

Migration in the Horn of Africa for $\delta = 0.73$

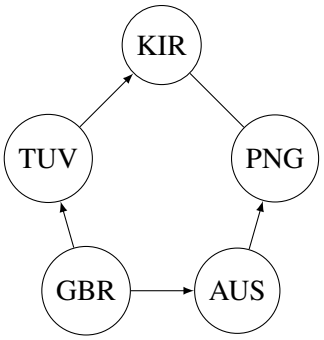

Migration in former British colonies, $\delta = 0.77$

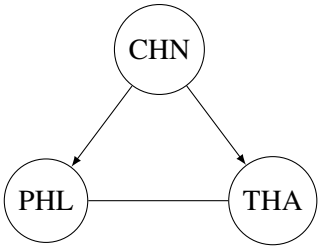

Migration in Southeast Asia, $\delta = 0.77$

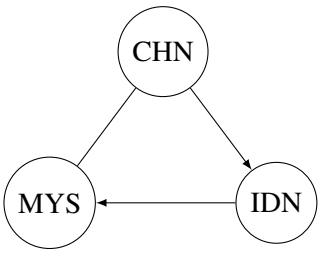

Migration in Southeast Asia, $\delta = 0.75$

FIGURE 23. **Top:** Two cycles corresponding to the left endpoints of the (Djibouti-Somalia-Uganda-Eritrea-Ethiopia) and (Kiribati-Papua New Guinea-Australia-United Kingdom-Tuvalu) persistence intervals listed in Table 5. **Bottom:** Two cycles corresponding to the left endpoints of the (China-Thailand-Philippines) and (China-Indonesia-Malaysia) persistence intervals listed in Table 5. **Meaning of arrows:** In each cycle, an arrow $s_i \to s_j$ means that $\omega_S(s_i, s_j) \leqslant \delta$, i.e. that s_j is a sink for the simplex $[s_i, s_j]$. We can verify separately that for $\delta = 0.77$, the Kiribati-Papua New Guinea simplex has the Solomon Islands as a δ-sink, and that the Philippines-Thailand simplex has Taiwan as a δ-sink. Similarly, the China-Malaysia simplex has Singapore as a δ-sink, for $\delta = 0.75$.

REFERENCES

[AA15] Michal Adamaszek and Henry Adams. The Vietoris-Rips complexes of a circle. *arXiv preprint arXiv:1503.03669*, 2015.

[AOTS15] Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *The American Economic Review*, 105(2):564–608, 2015.

[Atk72] Ronald H Atkin. From cohomology in physics to q-connectivity in social science. *International journal of man-machine studies*, 4(2):139–167, 1972.

[Atk75] Ron Atkin. *Mathematical structure in human affairs*. Crane, Russak, 1975.

[BBI01] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*, volume 33 of *AMS Graduate Studies in Math.* American Mathematical Society, 2001.

[BL14] Ulrich Bauer and Michael Lesnick. Induced matchings of barcodes and the algebraic stability of persistence. In *Proceedings of the thirtieth annual symposium on Computational geometry*, page 355. ACM, 2014.

[BO04] Albert-Laszlo Barabasi and Zoltan N Oltvai. Network biology: understanding the cell's functional organization. *Nature reviews genetics*, 5(2):101–113, 2004.

[Car09] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.

[CCSG+09a] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J Guibas, and Steve Y Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 237–246. ACM, 2009.

[CCSG+09b] Frédéric Chazal, David Cohen-Steiner, Leonidas J Guibas, Facundo Mémoli, and Steve Y Oudot. Gromov-hausdorff stable signatures for shapes using persistence. In *Computer Graphics Forum*, volume 28, pages 1393–1403. Wiley Online Library, 2009.

[CDS10] Gunnar Carlsson and Vin De Silva. Zigzag persistence. *Foundations of computational mathematics*, 10(4):367–405, 2010.

[CDSO14] Frédéric Chazal, Vin De Silva, and Steve Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1):193–214, 2014.

[CH13] CJ Carstens and KJ Horadam. Persistent homology of collaboration networks. *Mathematical Problems in Engineering*, 2013, 2013.

[CI08] Carina Curto and Vladimir Itskov. Cell groups reveal structure of stimulus space. *PLoS Computational Biology*, 4(10):e1000205, 2008.

[CM08] Gunnar Carlsson and Facundo Mémoli. Persistent clustering and a theorem of j. kleinberg. *arXiv preprint arXiv:0808.2241*, 2008.

[CM10] Gunnar Carlsson and Facundo Mémoli. Characterization, stability and convergence of hierarchical clustering methods. *Journal of Machine Learning Research*, 11:1425–1470, 2010.

[CM13] Gunnar Carlsson and Facundo Mémoli. Classifying clustering schemes. *Foundations of Computational Mathematics*, 13(2):221–252, 2013.

[CM15] Samir Chowdhury and Facundo Mémoli. Metric structures on networks and applications. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1470–1472, Sept 2015.

[CM16a] Samir Chowdhury and Facundo Mémoli. Distances between directed networks and applications. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6420–6424. IEEE, 2016.

[CM16b] Samir Chowdhury and Facundo Mémoli. Network motifs, distances between networks, and the stability of invariants. *preprint*, 2016.

[CM16c] Samir Chowdhury and Facundo Mémoli. Persistent homology of directed networks. *50th Asilomar Conference on Signals, Systems, and Computers*, page to appear, IEEE, 2016.

[CMRS13] Gunnar Carlsson, Facundo Mémoli, Alejandro Ribeiro, and Santiago Segarra. Axiomatic construction of hierarchical clustering in asymmetric networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 5219–5223. IEEE, 2013.

[CMRS14] Gunnar E. Carlsson, Facundo Mémoli, Alejandro Ribeiro, and Santiago Segarra. Hierarchical quasi-clustering methods for asymmetric networks. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014*, 2014.

[Cro78] Fred H Croom. *Basic Concepts of Algebraic Topology*. Springer-Verlag, New York, 1978.

[CSEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.

[CZCG04] Anne Collins, Afra Zomorodian, Gunnar Carlsson, and Leonidas J Guibas. A barcode shape descriptor for curve point cloud data. *Computers & Graphics*, 28(6):881–894, 2004.

[CZCG05] Gunnar Carlsson, Afra Zomorodian, Anne Collins, and Leonidas J Guibas. Persistence barcodes for shapes. *International Journal of Shape Modeling*, 11(02):149–187, 2005.

[DHR+16] Pawel Dlotko, Kathryn Hess, Levi Ran, Henry Markram, Eilif Muller, Max Nolte, Michael Reimann, and Martina Scolamiero. Topological analysis of the connectome of digital reconstructions of neural microcircuits. 2016.

[DMFC12] Yu Dabaghian, Facundo Mémoli, L Frank, and Gunnar Carlsson. A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Comput Biol*, 8(8):e1002581, 2012.

[Dow52] Clifford H Dowker. Homology groups of relations. *Annals of mathematics*, pages 84–95, 1952.

[DSC04] Vin De Silva and Gunnar Carlsson. Topological estimation using witness complexes. *Proc. Sympos. Point-Based Graphics*, pages 157–166, 2004.

[EGJ14] Matthew Elliott, Benjamin Golub, and Matthew O Jackson. Financial networks and contagion. *The American economic review*, 104(10):3115–3153, 2014.

[EH08] Herbert Edelsbrunner and John Harer. Persistent homology-a survey. *Contemporary mathematics*, 453:257–282, 2008.

[EH10] Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.

[EIK01] Alon Efrat, Alon Itai, and Matthew J Katz. Geometry helps in bottleneck matching and related problems. *Algorithmica*, 31(1):1–28, 2001.

[EJM15] Herbert Edelsbrunner, Grzegorz Jabłoński, and Marian Mrozek. The persistent homology of a self-map. *Foundations of Computational Mathematics*, 15(5):1213–1244, 2015.

[EK10] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010.

[ELZ02] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28(4):511–533, 2002.

[FL99] Patrizio Frosini and Claudia Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9(4):596–603, 1999.

[Fro92] Patrizio Frosini. Measuring shapes by size functions. In *Intelligent Robots and Computer Vision X: Algorithms and Techniques*, pages 122–133. International Society for Optics and Photonics, 1992.

[Ghr08] Robert Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.

[Ghr14] Robert Ghrist. *Elementary applied topology*. Createspace, 2014.

[GPCI15] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences*, 112(44):13455–13460, 2015.

[HMR09] Danijela Horak, Slobodan Maletić, and Milan Rajković. Persistent homology of complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(03):P03034, 2009.

[HNH+16] Yasuaki Hiraoka, Takenobu Nakamura, Akihiko Hirata, Emerson G Escolar, Kaname Matsue, and Yasumasa Nishiura. Hierarchical structures of amorphous solids characterized by persistent homology. *Proceedings of the National Academy of Sciences*, 113(26):7035–7040, 2016.

[HP06] Karen J Horowitz and Mark A Planting. Concepts and methods of the input-output accounts. 2006.

[HRS10] Daniel H Huson, Regula Rupp, and Celine Scornavacca. *Phylogenetic networks: concepts, algorithms and applications*. Cambridge University Press, 2010.

[Joh13] Jeffrey Johnson. *Hypernetworks in the science of complex systems*, volume 3. World Scientific, 2013.

[KKC+14] Arshi Khalid, Byung Sun Kim, Moo K Chung, Jong Chul Ye, and Daejong Jeon. Tracing the evolution of multi-scale functional networks in a mouse model of depression using persistent brain network homology. *NeuroImage*, 101:351–363, 2014.

[KNT10] Ravi Kumar, Jasmine Novak, and Andrew Tomkins. Structure and evolution of online social networks. In *Link mining: models, algorithms, and applications*, pages 337–357. Springer, 2010.

[KO97] Nigel J Kalton and Mikhail I Ostrovskii. Distances between banach spaces. *arXiv preprint math/9709211*, 1997.

[KPT07] Bala Krishnamoorthy, Scott Provan, and Alexander Tropsha. A topological characterization of protein structure. In *Data Mining in Biomedicine*, pages 431–455. Springer, 2007.

[LCK+11] Hyekyoung Lee, Moo K Chung, Hyejin Kang, Boong-Nyun Kim, and Dong Soo Lee. Computing the shape of brain networks using graph filtration and Gromov-Hausdorff metric. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2011*, pages 302–309. Springer, 2011.

[Lef42] Solomon Lefschetz. *Algebraic topology*, volume 27. American Mathematical Soc., 1942.

[Mas14] Anthony J Masys. *Networks and Network Analysis for Defence and Security*. Springer Science & Business Media, 2014.

[Mun84] James R Munkres. *Elements of algebraic topology*, volume 7. Addison-Wesley Reading, 1984.

[New03] Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.

[OD71] John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 34(1):171–175, 1971.

[ÖPSW11] Çaglar Özden, Christopher R Parsons, Maurice Schiff, and Terrie L Walmsley. Where on earth is everybody? the evolution of global bilateral migration 1960–2000. *The World Bank Economic Review*, 25(1):12–56, 2011.

[Pes14] Luiz Pessoa. Understanding brain networks and brain organization. *Physics of life reviews*, 11(3):400–435, 2014.

[PSDV13] Giovanni Petri, Martina Scolamiero, Irene Donato, and Francesco Vaccarino. Topological strata of weighted complex networks. *PloS one*, 8(6):e66506, 2013.

[Rob99] Vanessa Robins. Towards computing homology from finite approximations. In *Topology proceedings*, volume 24, pages 503–532, 1999.

[RS10] Mikail Rubinov and Olaf Sporns. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52(3):1059–1069, 2010.

[SCM16] Zane Smith, Samir Chowdhury, and Facundo Mémoli. Hierarchical representations of network data with optimal distortion bounds. *50th Asilomar Conference on Signals, Systems, and Computers*, page to appear, IEEE, 2016.

[SK04] Olaf Sporns and Rolf Kötter. Motifs in brain networks. *PLoS biology*, 2(11):e369, 2004.

[Spo11] Olaf Sporns. *Networks of the Brain*. MIT press, 2011.

[Spo12] Olaf Sporns. *Discovering the human connectome*. MIT press, 2012.

[TVJA] Andrew Tausz, Mikael Vejdemo-Johansson, and Henry Adams. Javaplex: A research software package for persistent (co) homology.

[Wei11] Shmuel Weinberger. What is... persistent homology? *Notices of the AMS*, 58(1):36–39, 2011.

[XW14] Kelin Xia and Guo-Wei Wei. Persistent homology analysis of protein structure, flexibility, and folding. *International journal for numerical methods in biomedical engineering*, 30(8):814–844, 2014.

[ZC05] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

E-mail address: chowdhury.57@osu.edu, memoli@math.osu.edu

(S. Chowdhury) DEPARTMENT OF MATHEMATICS, THE OHIO STATE UNIVERSITY. PHONE: (614) 292-6805.

(F. Mémoli) DEPARTMENT OF MATHEMATICS AND DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING, THE OHIO STATE UNIVERSITY. PHONE: (614) 292-4975, FAX: (614) 292-1479.